arXiv:1402.6118v1 [stat.ME] 25 Feb 2014

Approximate Models and Robust Decisions

James Watson and Chris Holmes*

February 26, 2014

Abstract

Statistical decisions based partly or solely on predictions from probabilistic models may be sensitive to model misspecification. Statisticians are taught from an early stage that "all models are wrong" but little formal guidance exists on how to assess the impact of model approximation, or how to proceed when optimal actions appear sensitive to model fidelity. This article presents one potential applied framework to address this. We discuss diagnostic techniques, including graphical approaches and summary statistics, to help highlight decisions made through minimised expected loss that are sensitive to model misspecification. We then derive formal methods for decision making under model misspecification by quantifying stability of optimal actions to perturbations within a neighbourhood of model space, defined via an information (Kullback-Leibler) divergence around the approximating model. This latter approach draws heavily from recent work in the robust control, macroeconomics and financial mathematics literature. We adopt a Bayesian approach throughout although the methods are agnostic to this position.

Keywords: Decision theory; Model Misspecification; *D-open* problems; Loss function; Kullback-Leibler divergence; Decision robustness.

1 Introduction

The central theme of the paper is this, that the consequence of statistical model misspecification is contextual and hence should be dealt with under a decision theoretic framework (Berger, 1985; Parmigiani & Inoue, 2009). As a trivial illustration consider the following: suppose that data arise from an exponential distribution, $x \sim \exp(\lambda)$, yet the statistician adopts a normal model, incorrectly assuming $x \sim N(\mu, \sigma^2)$. If interest is in the estimation of the mean $E[X]$ and the sample size is large there may be little consequence in the misspecification. However if the focus is on the probability of an interval event, say $X \in [a, b]$, then there might be far reaching consequences in using the model. Moreover, if we were seeking to estimate the mean, yet actions were critically sensitive to tail errors in mean specification then the normal model might no longer be adequate. Of course this is a toy problem, and careful model checking and refinement will help in reducing misspecification, but pragmatically, especially in modern high-dimensional data settings, it seems to us inappropriate to separate the issue of model misspecification from the consequences, context and rationale of the modelling exercise.

*Chris Holmes is Professor of Statistics, Department of Statistics, University of Oxford, UK. email: cholmes@stats.ox.ac.uk

Statisticians are taught from an early stage that "essentially, all models are wrong, but some are useful" (Box & Draper, 1987). By "wrong" we will take to mean misspecified and by "useful" we will take to mean helpful for aiding actions (taking decisions). We will refer to such situations as *D-open* problems, to highlight that Nature's true model is outside of the decision makers knowledge domain, c.f. *M-open* in Bayesian statistics which refers to problems in updating beliefs when the model is known to be wrong (Bernardo & Smith, 1994). We will adopt a Bayesian standpoint throughout although the approach we develop is generic. We will assume there is uncertainty in some aspect of the world[1], $\theta \in \Theta$, which if known would determine the loss in taking an action a as quantified through a real-valued measurable loss-function, $L_a(\theta)$. The loss will often be a joint function of states and observables, $L_a(\theta, x)$, although we shall suppress this notation for convenience. Uncertainty in θ is characterised via a probability distribution $\pi_I(\theta)$ given all available information I. Without loss of generality we will assume that θ relates to parameters of a probability model and information I is in the form of data, $x \in \mathcal{X}$, and a joint model $\pi(x, \theta)$, such that,

$$\pi_I(\theta) \equiv \pi(\theta|x) \propto f(x;\theta)\pi(\theta),$$

where $\pi(x, \theta)$ is factorised according to the sampling distribution (or likelihood) $f(x;\theta)$ and the prior $\pi(\theta)$; although more generally $\pi_I(\theta)$ simply represents the statisticians best current beliefs about the value of the unknown state θ. Following the axiomatic approach of Savage (1954) the rational coherent approach to decision making is to select an action $\hat{a}$ from the set of available actions $a \in A$ so as to minimise expected loss,

$$\hat{a} = \arg\min_{a \in A} E_{\pi_I(\theta)}[L_a(\theta)]. \tag{1}$$

This underpins the foundations of Bayesian statistics (Bernardo & Smith, 1994). The problem is that (1) assumes perfect precision in specifying $\pi(x, \theta)$. In reality the model $\pi(x, \theta)$ is misspecified, such that the decision maker acknowledges that $f(x;\theta)$ may not be Nature's true sampling distribution or $\pi(\theta)$ does not reflect all aspects of prior subjective beliefs in $f(x;\theta)$ or on the marginal $\pi(x) = \int_\theta \pi(x, \theta)d\theta$. This paper presents diagnostics and formal methods to assist in exploring the potential impact of this misspecification.

It is important to note that we will not spend much time on the area of pure inference problems, such as robust estimation of summary functionals, for which there is a substantial literature (Huber, 2011) or on recent work on the use of loss functions to construct posterior models (Bissiri & Walker, 2012b; Bissiri *et al.*, 2013). We shall also pass quickly over the use of conventional prior sensitivity analysis and robust "heavy tailed" priors and likelihoods. We are principally concerned with ex-post settings where $\pi(x, \theta)$ has been specified to the best of the modellers ability under the practical constraints of computation and time, and where concerns arise as to whether $\pi(x, \theta)$ models Nature's truth $\pi(x)$ to sufficient precision. This is particularly important when θ pertains to a high-dimensional complex model or to the value of a future predicted observation.

There is a rich literature in Bayesian statistics on model robustness, the vast majority of which relates to sensitivity to specification of the prior $\pi(\theta)$. We review the material in detail below but mention here the overviews in Berger (1994), Rios Insua & Ruggeri (2000) and Ruggeri *et al.* (2005). Bayesian robustness was a highly active area through the 1980s to mid-1990s. Interest tailored off somewhat since that time, principally due to the arrival of computational methods such as Markov chain Monte Carlo (MCMC) coupled with hierarchical models, non-linear models and nonparametric priors, see e.g. Chipman *et al.* (1998), Robert & Casella (2004), Rasmussen & Williams (2006), Denison *et al.* (2002), and Hjort *et al.* (2010). These

[1]Savage (1954) refers to Θ as the "small world" relevant to the decision.

methods allow for very flexible model specifications alleviating the historic concern that $\pi(x, \theta)$ was indexing a restrictive sub-class of models. However, a number of recent factors merit a reappraisal. In the 1990s and 2000s computational advances and hierarchical models broadly outpaced the complexity of data sets being considered by statisticians. In more recent times very high-dimensional data are becoming common, the so called "big data" era, whose size and complexities prohibit application of fully specified carefully crafted models, (e.g. (National Research Council: *et al.*, 2013), Chapter 7). Related to this, approximate probabilistic inference techniques that are misspecified by design have emerged as important tools for applied statisticians tackling complex inference problems. For example, models involving composite likelihoods, integrated nested Laplace approximations (INLA), Variational Bayes, Approximate Bayesian Computation (ABC), all start with the premise of misspecification, see e.g. Beaumont *et al.* (2002), Fearnhead & Prangle (2012), Marjoram *et al.* (2003), Marin *et al.* (2012), Minka (2001), Ratmann *et al.* (2009), Rue *et al.* (2009), Varin *et al.* (2011), and Wainwright & Jordan (2003). Finally there have been recent developments in coherent risk measures within the macroeconomics and mathematical finance literature, building on areas of robust control, which are of importance and relevance to statisticians, as outlined in Section 2 below.

The rest is as follows. In Section 2 we review some background literature on decision robustness and quantification of expected loss under model misspecification. In Section 3 we then propose diagnostic tools to assist applied statisticians in identifying actions which may be sensitive to model fidelity. Section 4 presents formal methods for summarising decision stability, by exploring the consequence of misspecification within local neighbourhoods around the approximate model. Section 5 contains illustrations. Conclusions are made in Section 6.

2 Background on decisions under model misspecification

We first review some of the background literature on decisions made under model misspecification.

2.1 Minimax

The first axiomatic approach to robust statistical decision making was made by Wald (1950). In the absence of a true model, Wald interpreted the decision problem as a zero sum two-person game, following Von Neumann and Morgenstern's work on game theory (Von Neumann & Morgenstern, 1947). To be robust the statistician protects himself against the worst possible outcome, selecting an action $\hat{a}$ according to the minimax rule, which for the purposes of this paper we can consider as[2],

$$\hat{a} = \arg\min_{a \in A} \left[\max_{\theta \in \Theta} L_a(\theta) \right].$$

This is akin to the decision maker playing a two-person game with a malevolent Nature, where losses made by one agent will be gained by the other (zero sum). On selection of an action, Nature will select the worst possible outcome, equivalent to the assumption of a point mass distribution taken reactively to your choice of action,

$$\delta_{\theta_a^*}(\theta),$$

[2]Wald's original work considered selection of decision functions, $\delta(x) \in A$, by non-conditional loss quantified as frequentist risk, $R[F_X, \delta(x)] = \int L(\delta, x) F(dx)$, with $x \in \mathcal{X}$ from unknown distribution F_X.

where,

$$\theta_a^* = \arg\max_{\theta \in \Theta} L_a(\theta).$$

Although elegant in its derivation the minimax rule has severe problems from an applied perspective. The decision maker following the minimax rule is not rational and treats all situations with extreme pessimism. It assumes that Nature is reactive in selecting $\delta_{\theta_a^*}(\theta)$ for your choice of $a \in A$ irrespective of the evidence from existing information I on the plausible values of θ. Subsequent to Wald there has been considerable work to develop more applied procedures that protect against less extreme outcomes.

2.2 Robust Bayesian statistics

Under a strict Bayesian position there is no issue with model robustness. You precisely specify your subjective beliefs through $\pi(x, \theta)$ and condition on data to obtain posterior beliefs, taking actions according to the Savage axioms. However, even the modern founders of Bayesian statistics acknowledged issues with an approach that assumes infinite subjective precision,

> "Subjectivists should feel obligated to recognise that any opinion (so much more the initial one) is only vaguely acceptable... So it is important not only to know the exact answer for an exactly specified initial problem, but what happens changing in a reasonable neighbourhood the assumed initial opinion." De Finetti, as quoted in Dempster (1975)

> "...in practice the theory of personal probability is supposed to be an idealization of one's own standard of behaviour; that the idealization is often imperfect in such a way that an aura of vagueness is attached to many judgements of personal probability..." Savage (1954)

As Berger points out, many people somewhat distrust the Bayesian approach as "Prior distributions can never be quantified or elicited exactly (i.e. without error), especially in finite amount of time" – Assumption II in Berger (1984). In which case what does the resulting posterior distribution $\pi(\theta|x)$ actually represent?

An intuitive solution is to first specify an operational model π_0, to the best of your available time and ability, and then investigate sensitivity of inference or decisions to departures around π_0, typically assuming that $f(x;\theta)$ is known so that divergence is with respect to the prior. This idea has origins in the work of Robbins (1952) and Good (1952) with many important contributions since that time. We mention just a few pertinent areas below, referring the interested reader to the review articles of Berger (1984), Berger (1994), Wasserman (1992), and Ruggeri *et al.* (2005), as well as the collection of papers in the edited volumes of Kadane (1984) and Rios Insua & Ruggeri (2000).

Robust Bayesian methods are usefully classified as either "local" or "global". Local approaches look at functional derivatives of posterior quantities of interest with respect to perturbations around the baseline model, e.g. Ruggeri & Wasserman (1993) Sivaganesan (2000); see also Kadane & Chuang (1978) who consider asymptotic stability of decision risk. Global approaches consider variation in a posterior functional of interest, $\psi = \int h(\theta)\pi(\theta|x)d\theta$, within a neighbourhood $\pi \in \Gamma$ centred around the prior model π_0. A typical quantity would be the range $(\psi_{\inf}, \psi_{\sup})$ where $\psi_{\inf} = \inf_{\pi\in\Gamma} \psi$ and $\psi_{\sup} = \sup_{\pi\in\Gamma} \psi$. The challenge is to define the nature and size of Γ so as to capture plausible ambiguity in π_0, while taking into account factors such as ease of specification and computational tractability, Berger (1994; 1985 section 4.7). One

important example is the ϵ-*contamination* neighbourhood (Berger & Berliner, 1986) formed by the mixture model,

$$\Gamma = \{\pi = (1-\epsilon)\pi_0 + \epsilon q, q \in \mathcal{Q}\},$$

where ϵ is the perceived contamination error in π_0 and $\mathcal{Q}$ is a class of contaminant distributions. It is usual to restrict $\mathcal{Q}$ so that it is not "too big", for instance by including only uni-modal distributions Berger (1994), for which it is shown that the solutions have tractable form. Other approaches consider frequentist risk, such as Γ-*minimax* that investigates the minimax Bayes (frequentist) risk of ψ_{sup} for $\pi \in \Gamma$ whereas *conditional* Γ-*minimax* procedures (Vidakovic, 2000) study the maximum expected loss across prior distributions within Γ, this being perhaps closest to the approach we develop here.

A key distinction between these approaches and that developed here, is that we shall be concerned with robustness to misspecification of the posterior marginal distribution on only those states θ that enter into the loss function $L_a(\theta)$. This facilitates application to high-dimesional problems for which specification of Γ may be difficult (Sivaganesan, 1994) and tackles the thorny issue that changing the likelihood changes the interpretation of the prior (Ruggeri *et al.* (2005), page 635).

2.3 Robust control, macroeconomics and finance

Independent of the above developments in statistics, control theorists were investigating robustness to modelling assumptions. Control theory broadly concerns optimal intervention strategies (actions) on stochastic systems so as to maintain the process within a stable regime. Hence it is not surprising that decision stability is an important issue. When the system is linear with additive normal (white) noise the optimal intervention is well known Whittle (1990). Robust control theory, principally developed by Whittle, considers the case when Nature is acting against the operator through stochastic buffering by non-independent noise, see Whittle (1990). Whittle established that under a malevolent Nature with a bounded variance an optimal intervention can be calculated using standard recursive algorithms.

In Economics one early criticism of the Savage axioms was that the framework could not distinguish between different types of uncertainty. Gilboa & Schmeidler (1989) developed a theory of maxmin Expected Utility in part to counter the famous Ellsberg paradox[3]which extends standard Bayesian inference to a setting with multiple priors in the form of a closed convex set Γ. An action is then scored by its expected loss under the least favourable prior within that set. Their 1989 paper formalises this and provides a solution to the Ellsberg paradox. When Γ contains only one prior, we are back again in the usual Bayesian setting. The set Γ can be seen as describing the decision-maker's aversion to uncertainty. This work is closely related to Γ-*minimax* (for which the Ellsberg paradox is also used as a motivating example, see section 1 of Vidakovic (2000)).

Again working in economics, Hansen and Sargent in a series of influential papers (e.g. 2001a, 2001b), generalised ideas from Whittle (1990) and Gilboa & Schmeidler (1989) motivated by problems in macroeconomic time series. They define a robust action as a local-minimax act

[3]The standard setting for the Ellsberg paradox is as follows: imagine two urns each containing 100 balls and every ball is either red or blue. One is told that the first urn (A) has 50 red balls and 50 blue balls exactly. No more information is given about the second urn (B). Suppose you win 100$ if you pick a red ball, which urn would you choose? So there exists a set of alternatives which are equal in expected value (under any reasonable prior) but which appear to have different empirical preferences.

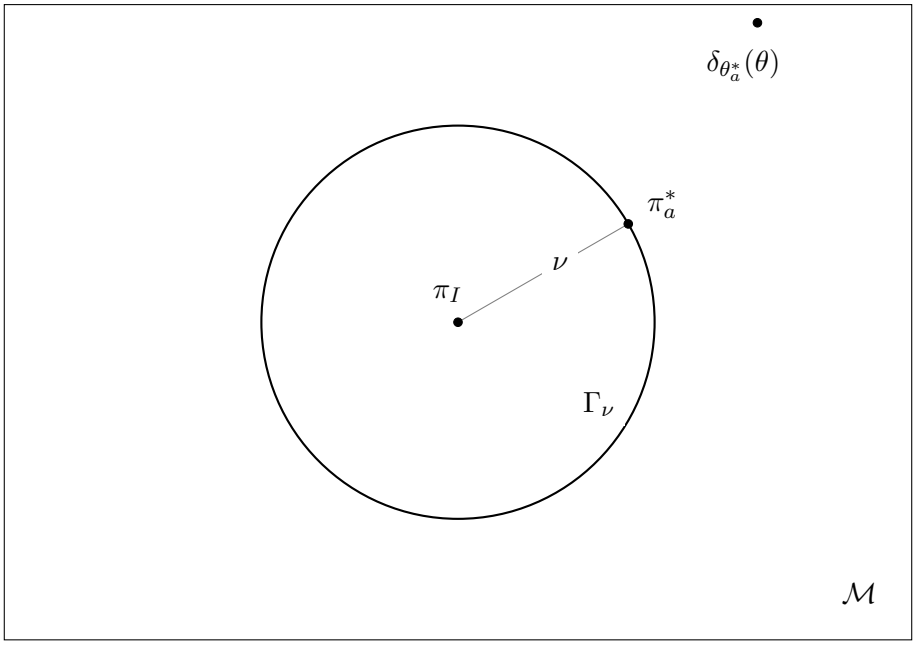

Figure 1: Graphical representation of local-minimax model π_a^* within a Kullback-Leibler ball of radius ν around the reference model π_I, with global (Wald's) minimax density $\delta_{\theta_a^*}(\theta)$.

within a Kullback-Leibler (KL) neighbourhood of $\pi_I(\theta)$ through exploration of,

$$\psi_{\text{sup}}^{(a)} = \sup_{\pi \in \Gamma} E_\pi[L_a(\theta)]$$

where Γ denotes a KL ball around π_I,

$$\Gamma = \{\pi : \int \pi(\theta) \log\left(\frac{\pi(\theta)}{\pi_I(\theta)}\right) d\theta \leq C\}.$$

We will use π_a^* or π_{sup}^* interchangeably to denote the corresponding local-minimax distribution,

$$\pi_a^* = \arg\sup_{\pi \in \Gamma} E_\pi[L_a(\theta)].$$

Figure 1 shows a pictorial representation of this constrained minimax rule, where the reference distribution π_I is a point in the space of distributions $\mathcal{M}$ (represented by the rectangle) and the least favourable distribution π_a^* is contained within the neighbourhood Γ_ν (represented by the circle of radius ν). The Wald minimax distribution is given by $\delta_{\theta_a^*}(\theta)$. Hansen and Sargent showed how π_a^* and $\psi_{\text{sup}}^{(a)}$ can be computed for dynamic linear systems with normal noise, see Hansen & Sargent (2008) for a thorough review and references.

Breuer & Csiszár (2013a, 2013b), building on the work of Hansen and Sargent, derived corresponding results for arbitrary probability measures $\pi_I(\theta)$. Under mild regularity conditions, and using results from exponential families and large deviation theory they obtain the exact form of π_a^* for any $\pi_I(\theta)$ given the KL ball of size C, as well as an estimate for $\psi_{\text{sup}}^{(a)}$, see also Ahmadi-Javid (2011, 2012). In Section 4 we derive the same result using an alternative, less general, but perhaps more intuitive proof. Before considering these formal methods we shall start with exploratory diagnostics and visualisation methods.

3 *D-open* diagnostics

All good statistical data analysis begins with graphical exploration of pertinent features and evaluation of summary statistics before formal modelling takes place. In this section we consider graphical displays to aid in understanding when and where actions are sensitive to modelling assumptions. Despite the importance of graphical statistics, there are few if any established tools for investigation of decision stability, in contrast to the multitude of methods for investigating model discrepancy and misspecification (Belsley *et al.*, 2005) (Gelman, 2007).

Throughout, we consider that $\pi_I(\theta)$ can be approximated by a bag of Monte Carlo samples $\{\theta_i\}_{i=1}^m$ with $\theta_i \sim \pi_I(\theta)$ and that $\pi_I(\theta_i)$ can be evaluated point wise up to a normalizing constant. We shall also assume that the loss is bounded and that after linear transformation $\inf_{a\in A,\theta\in\Theta}[L_a(\theta)] = 0$.

3.1 Value at Risk (Quantile-Loss)

A primary tool for assessing the sensitivity with respect to misspecification of a functional of interest is the distribution of loss, where Z_a denotes the random loss variable under $\pi_I(\theta)$:

$$F_{Z_a}(z) = Pr(Z_a \leq z) = \int_{\theta\in\Theta} I[L_a(\theta) \leq z]\pi_I(d\theta),$$

where $I[\cdot]$ is the identity function. We use notation $f_a(z) = F_{Z_a}(dz)$ to denote the corresponding density function and $F_{Z_a}^{-1}(q)$, for $q \in [0,1]$, is the inverse cumulative distribution or quantile function. For a given $q \in [0,1]$ it is possible to characterise the value of an action a by its quantile loss or "value at risk" (VaR; terminology used in finance) $F_{Z_a}^{-1}(1-q)$. Rostek (2010) developed an axiomatic framework in which decision-makers can be uniquely characterised by a quantile $1-q$ and rational behaviour (optimal action) is defined as choosing $\hat{a} := \arg\min_{a\in\mathcal{A}} F_{Z_a}^{-1}(1-q)$. Note that this incorporates the minimax rule by taking $q = 0$. The author argues that quantile maximisation is attractive to practitioners as its key characteristic is robustness, specifically to misspecification in the tails of the loss distribution. Although single quantiles discard much information contained in $[\pi_I(\theta), L_a(\theta)]$, plotting this function allows for immediate visualisation of how much of the tails are taken up by high loss (low utility) events. With a bag of samples $\theta_1, .., \theta_m \sim \pi_I(\theta)$, this is easily approximated:

1. Sort the realised loss values, $z_i^{(a)} = L_a(\theta_i)$, from lowest to highest $\{z_{\upsilon_a(1)}^{(a)} \leq z_{\upsilon_a(2)}^{(a)} \leq \ldots \leq z_{\upsilon(m)}^{(a)}\}$, where $\upsilon_a(\cdot)$ defines the sort mapping.
2. Approximate $F_{Z_a}^{-1}(1-q)$ by linear interpolation of the points $\left(x = \frac{k}{m}, y = z_{\upsilon_a(k)}^{(a)}\right)$.

To give a simple illustration of these plots, we consider three alternatives with the following loss distributions (all over the interval $[0,1]$): a_1 with uniform loss on $[0,1]$, a_2 with a Beta$(5,2)$ distribution and a_3 with a Beta$(5,2)$ distribution. We characterise a_1 as a *stable decision*, maximising the entropy over the set of possible losses. a_2 is illustrative of a decision with high upside (i.e. possibility of achieving low loss but most of the probability mass is on high loss, near 1), and a_3 a decision with high downside. Figure 2 shows the loss distributions and VAR plots (inverse loss) for these three decisions, highlighting the differences in possible shapes of the VAR curve for stable, high upside and high downside decisions.

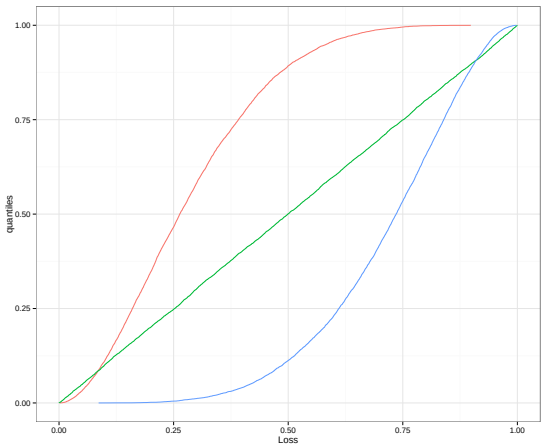

(a) Loss distribution for example given in section 3.1

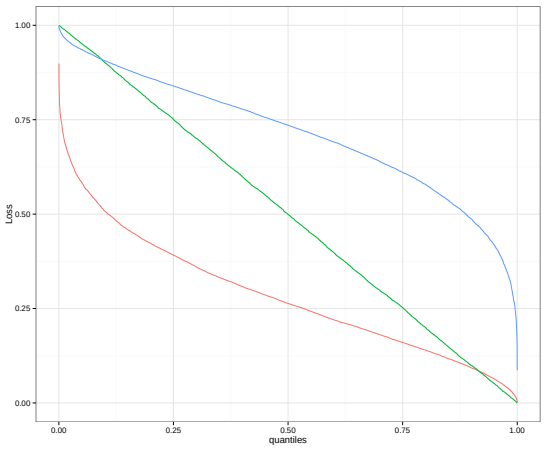

(b) Illustrative plot: VaR plot as explained in section 3.1

Figure 2: Detecting sensitivity in decisions: Loss distribution and inverse loss distribution.

3.2 Conditional value at risk (upper-trimmed mean utility)

In finance the conditional value at risk (CVaR) (Rockafellar & Uryasev, 2000) is another popular alternative to expected loss (or utility). To a statistician it represents the lower trimmed mean of loss (or upper trimmed mean of utility),

$$G_a(q) = \frac{1}{q}\int_{F_{Z_a}^{-1}(1-q)}^{\infty} z f_a(z) dz.$$

This gives the expected value of an action conditional on the event (θ) occurring above a quantile of loss (lowest of utility). q can be seen as regulating the amount of pessimism towards Nature, with $\lim_{q\to 0} \sup_a G_a(q)$ corresponding to the minimax rule.

Another strategy for taking in to account model misspecification is by considering the two-sided *trimmed expected loss*, defined as:

$$H(q) = \frac{1}{1-q}\int_{F_{Z_a}^{-1}(q/2)}^{F_{Z_a}^{-1}(1-q/2)} z f_a(z) dz$$

This is a robust measure of expected loss formed by discarding events with highest and lowest loss. Both these statistics are easily approximated using the bag of samples $\{\theta_i\}_{i=1}^n$ and the sort mapping υ defined previously. We use the linear interpolation,

$$\hat{G}_a(\frac{k}{m}) = \frac{1}{k}\sum_{i=1}^{k} L_a(\theta_{\upsilon(i)}), \quad k = 0, .., m$$

For a set of actions $\mathcal{A}$, it is possible to quantify the stability of the optimal action a^* evaluated under expected loss, by observing the first CVaR crossing point. That is to say the first value $q \in [0, 1]$ such that a^* is no longer optimal, evaluated under CVaR(q).

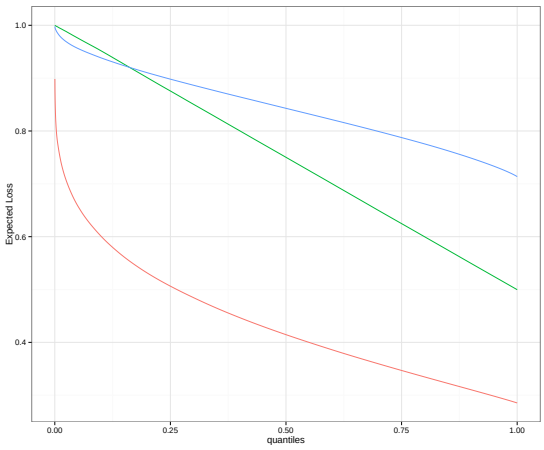

(a) Illustrative plot: CVaR plot as explained in section 3.2

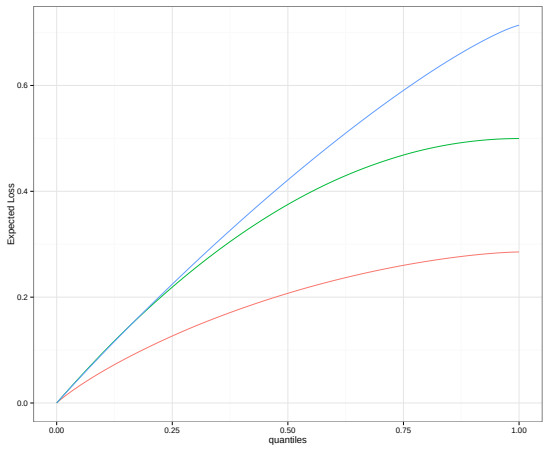

(b) Illustrative plot: CER plot as explained in section 3.3

Figure 3: Sensitivity to misspecification: Conditional Value at Risk and Cumulative Expected Return.

3.3 Cumulative Expected Return (CER)

The *Cumulative Expected Return* (CER) function for action a, defined as,

$$J_a(q) = \int_{F_{Z_a}^{-1}(1-q)}^{\infty} z f_a(z) dz.$$

for $q \in [0, 1]$. The CER-plot is a monotone decreasing function $J_a(q)$ and an informative graph for highlighting decision sensitivity. As shown in Figure 3b, the overall shape of $J_a(q)$ provides a qualitative description of decision sensitivity. An action with CER-plot that is steeply rising as $q \to 0$ is 'heavily downside', with expected-loss driven by low-probability high loss outcomes, while CER-plot rising at 1 indicates 'heavy upside'. In particular we note,

- $J_a(q)$ quantifies the contribution to the expected loss of action a, from the $100 \times (1-q)\%$ set of outcomes with greatest loss.
- $J_a(1) = E_{\pi_I(\theta)}[L_a(\theta)]$, is the expected loss of action a, and $\hat{a} = \arg\max_{a \in A} J_a(1)$ is the optimal Savage action.
- $J'_a(q) = \inf_{z^* \in \mathbb{R}^+} \{Pr(Z_a \leq z^*) = 1 - q\}$, the gradient of the curve at $J_a(q)$ gives the threshold loss value z^*, such that under action a we can expect with probability $(1-q)$ the outcome to have loss less than or equal to z^*.This is the "value-at-risk" of action a outlined above, e.g. Pritsker (1997).
- $J'_a(0) = \sup_{\theta \in \Theta} L_a(\theta)$, and hence the Wald minimax action is given by: $\tilde{a} = \arg\min_{a \in A} J'_a(0)$ (the action with steepest gradient as $q \to 0$).

3.4 Loss-Density Plots

It is often the case that rare events carry greatest impact. It is therefore informative to explore the relationship between $[\pi_I(\theta), L_a(\theta)]$. Given a bag of samples $\{\theta_i\}_{i=1}^m$ and the corresponding realised loss values $z_i^{(a)}$, we can plot the relationship between the density at the points θ_i, $\pi_I(\theta_i)$ (up to a normalising constant) and the loss $z_i^{(a)}$. This can highlight whether or not there is a

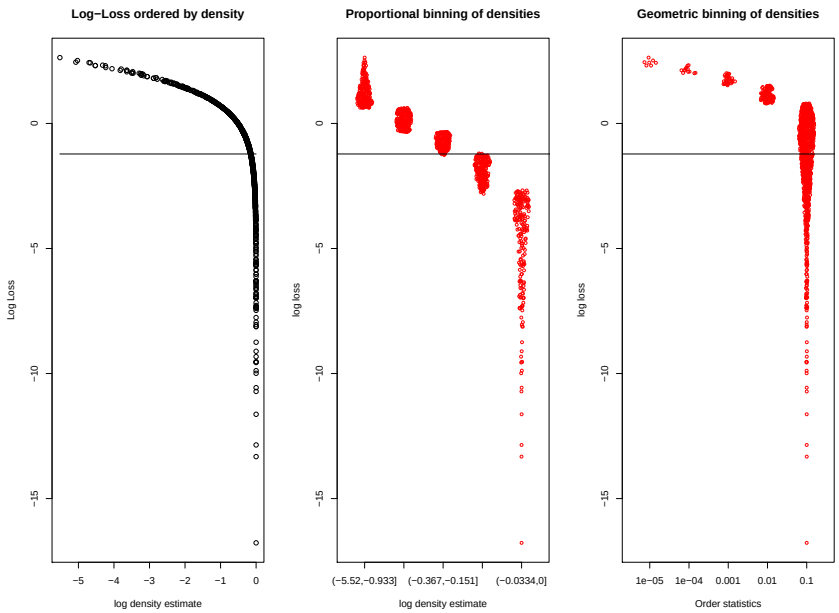

Figure 4: Log Loss-Density plot: Normal fitted to Student-t

clear dependence between the tails of $\pi_I(\theta)$ and high loss. As an example we look at generating data from a student-t distribution ($t = 4$), and then fitting a normal distribution with reference priors. Figure 4 shows this, with three plots, corresponding to no binning, proportional binning (uniform intervals) and geometric binning (logarithmically uniform intervals) of loss.

If the estimates $\pi_I(\theta_i)$ are binned either proportionally (into k bins of equal size across the interval $[\min_i\{\pi_I(\theta_i)\}, \max_i\{\pi_I(\theta_i)\}]$) or logarithmically (into k bins with sizes proportional to 10^i, for $i = 1, .., k$ across the same interval), then outliers can be spotted in the distribution over loss induced by each bin. As a diagnostic plot this could highlight under-explored areas of loss, e.g. for (relatively) large interval of density in $\pi_I(\theta)$, a set of outliers might correspond to a mode in the distribution that is under-represented. It also clearly illustrates in this example that the expected loss is heavily influenced by low probability events.

4 *D-open* formal methods

To formally investigate the stability of decisions to model misspecification we propose to follow the approach taken in Hansen & Sargent (2001b) and study the variation of expected loss ψ, over models within a KL ball Γ around the marginal posterior density, $\pi_I(\theta)$, of the approximating model on the states that enter into the loss function. We will assume after linear transformation that the loss can be bounded, a reasonable assumption for almost all applied problems. It is well known that the KL divergence is not symmetric, $KL(\pi^* \parallel \pi) \neq KL(\pi \parallel \pi^*)$ for $\pi^* \neq \pi$, and we shall consider each neighbourhood in turn. Interestingly, each poses different computational and analytical tractability issues.

4.1 $\mathbf{KL}(\pi \parallel \pi_I)$

We first consider the set of models where the divergence of π_I from every model is below some threshold. This might be considered the more natural setting as under this situation the KL divergence represents the expected self-information log-loss in using an approximate model

π_I when Nature is providing outcomes according to the probability law π. Surprisingly this situation leads to a local-minimax solution with simple form.

Theorem 4.1. *Let* $\pi^*_{\sup} = \arg\sup_{\pi\in\Gamma} E_\pi[L_a(\theta)]$, *with* $\Gamma = \{\pi : KL(\pi \parallel \pi_I) \leq C\}$ *for* $C \geq 0$. *Then the solution* $\pi^*_{\sup}$ *is unique and has the following form,*

$$\pi^*_{\sup} = Z^{-1}\pi_I(\theta)\exp[h(C)L_a(\theta)] \tag{2}$$

where $Z_C = \int \pi_I(\theta)\exp[h(C)L_a(\theta)]d\theta$ *is the normalising constant or partition function, for which we assume* $Z_C < \infty$, *and* $h(C)$ *is a non-negative real valued monotone function.*

Proof. The function minimisation problem, $\pi^*_{\sup} = \arg\max_{\pi\in\Gamma} E_\pi[L_a(\theta)]$, has a Lagrange dual form, see for example Hansen *et al.* (2006) (pages 58-60),

$$\pi^*_{\sup} = \arg\inf_{\pi\in\Gamma}\left\{E_\pi[-L_a(\theta)] + \lambda^{-1}KL(\pi \parallel \pi_I)\right\}$$

for some $\lambda \in [0,\infty)$, with $C < C' \implies \lambda < \lambda'$. Hence,

$$\begin{aligned}
\pi^*_{\sup} &= \arg\inf_{\pi\in\Gamma}\left\{\int -L_a(\theta)\pi(\theta)d\theta + \lambda^{-1}\int \pi(\theta)\log\left(\frac{\pi(\theta)}{\pi_I(\theta)}\right)d\theta\right\} \\
&= \arg\inf_{\pi\in\Gamma}\left\{\int \pi(\theta)\log\left(\frac{\pi(\theta)}{\pi_I(\theta)\exp[\lambda L_a(\theta)]}\right)d\theta\right\} \\
&\propto \pi_I(\theta)\exp[\lambda L_a(\theta)] \qquad (3)
\end{aligned}$$

The uniqueness arises from the convexity of the KL loss. The result follows. □

By a similar argument the distribution of minimum expected loss follows $\pi^*_{\inf} \propto \pi_I(\theta)\exp[-\lambda L_a(\theta)]$. Note by assuming bounded loss functions we can ensure the integrability of the densities by taking θ to have bounded support. Breuer & Csiszár (2013a) and Ahmadi-Javid (2012) derive the same result more generally but perhaps less intuitively. Breuer & Csiszár (2013a) gives more general conditions on when the solution exists.

The Γ-constrained minimax distributions, $\{\pi^*_{\inf}, \pi^*_{\sup}\}$, have an interpretable form as exponentially tilted densities, tilted toward the exponentiated loss function, with weighting λ a monotone function of the neighbourhood size C. For linear loss, $L_a(\theta) = A\theta$, the constrained-minimax $\pi^*_{\sup}$ is the well known Esscher Transform used for option pricing in actuarial science.

Using results in Bissiri *et al.* (2013), Bissiri & Walker (2012a) we are are able to state the following result regarding the uniqueness of Kullback-Leibler divergence for the neighbourhood set criteria.

Theorem 4.2. *Let* $\pi^*_{\sup}(x, \pi_I)$ *be the solution obtained by*

$$\pi^*_{\sup}(x, \pi_I) = \arg\inf_{\pi\in\Gamma}\left\{E_\pi[-L_a(\theta)] + \lambda^{-1}D(\pi \parallel \pi_I)\right\}$$

with data $x = \sum_{i=1}^n x_i$, *centering distribution* π_I, *and arbitrary divergence measure* $D(\cdot \parallel \cdot)$. *Moreover, let* x *be partitioned as* $x = \{x_1, x_2\}$, *for* $x_1 = \sum_{i=1}^{m-1} x_i$, $x_2 = \sum_{i=m}^n x_i$, $m \in \{0,\ldots,n-1\}$. *For coherence we require,*

$$\pi^*_{\sup}(x, \pi_I) \equiv \pi^*_{\sup}(x_2, \pi^*_{\sup}(x_1, \pi_I))$$

for any partition of x, *including* $x_2 = \emptyset$. *That is, the solution using a partial update involving* x_1, *which is subsequently updated to include* x_2, *should coincide with the solution obtained using all of the data, for any partition. Then for coherence* $D(\cdot \parallel \cdot)$ *is the Kullback-Leibler divergence.*

Proof, see the Appendix here, reproduced from Bissiri *et al.* (2013). □

This theorem shows that the KL is the only divergence measure to provide coherent updating of the local-minimax distribution.

Conditional Γ-minimax priors. When the loss function contains all of the parameters in the model this result allows us to consider the local-minimax operation as if applied initially to the prior $\pi(\theta)$,

$$\pi^*_{\text{sup}}(\theta) \propto e^{\lambda L_a(\theta)}\pi(\theta)$$

and then the prior can be updated to the posterior local-minimax solution by Bayes Theorem

$$\pi^*_{\text{sup}}(\theta|x) \propto f(x;\theta)\pi^*_{\text{sup}}(\theta)$$

where $f(x;\theta)$ is the likelihood. Note that the KL divergence is the only divergence to ensure this coherency, and also that the "prior" $\pi^*_{\text{sup}}(\theta)$ is data dependent if $L_a(\theta)$ is a function of x.

Following Theorem 4.1, the corresponding range of expected losses for each action $(\psi^{(a)}_{\text{inf}}, \psi^{(a)}_{\text{sup}})$ can then be plotted as a function of C for each potential action. Formally we should write $[\psi^{(a)}_{\text{inf}}(C), \psi^{(a)}_{\text{sup}}(C)]$ although for ease of notation we will often suppress C from the expression unless clarity dictates. The constraint $KL(\pi \parallel \pi_I) \leq C$ will result in $KL(\pi^*_{\text{sup}} \parallel \pi_I) = C$ as the expected loss can always be increased by diverging toward $\delta_{\theta^*}(\theta)$ for any distribution with $KL(\pi \parallel \pi_I) < C$. Substituting the solution (2) into the KL divergence function gives,

$$\begin{aligned} KL(\pi^*_{\text{sup}} \parallel \pi_I) &= \int \pi^*_{\text{sup}}(\theta) \log\left(Z^{-1}\exp[\lambda L_a(\theta)]\right) d\theta \\ &= \lambda E_{\pi^*_{\text{sup}}}[L_a(\theta)] - \log Z_\lambda \end{aligned}$$

So, given neighbourhood size λ, the KL divergence $KL(\pi^*_{\text{sup}} \parallel \pi_I)$ is λ times the expected loss under π^*_{sup} minus the log partition function. Moreover, by Jensen's inequality,

$$\begin{aligned} KL(\pi^*_{\text{sup}} \parallel \pi_I) &= \lambda E_{\pi^*_{\text{sup}}}[L_a(\theta)] - \log E_{\pi_I}[\exp(\lambda L_a(\theta)] \\ &\leq \lambda\left[E_{\pi^*_{\text{sup}}}[L_a(\theta)] - E_{\pi_I}[L_a(\theta)]\right] \end{aligned}$$

The KL divergence is bounded above by λ times the difference in expected loss between the approximating and the contained minimax models.

By plotting out the interval $[\psi^{(a)}_{\text{inf}}(C), \psi^{(a)}_{\text{sup}}(C)]$ for each action as a function of KL divergence constraint $C : KL(\pi \parallel \pi_I) \leq C$ we can look for crossing points between the supremum loss $\psi^{(\hat{a})}_{\text{sup}}$ under the optimal action $\hat{a}$ chosen by the approximating model and the infimum loss under all other actions, $\psi^{(a')}_{\text{inf}} = \inf_{a'}\{\psi^{(a')}_{\text{inf}} : a \in A \setminus \hat{a}\}$

In a classical setting, the notion of *admissibility* helps define a smaller class of actions that can then be further scrutinized in order to choose an optimal decision. A decision is denoted inadmissible if there does not exist a θ such that its risk function (frequentist) is minimal (with respect to the other decisions) at θ. We note that in a Bayesian context, because the expected loss is a single quantity used to classify actions, only the action which minimizes expected loss is admissible. However if we consider the set of posterior distributions contained within a Kullback-Leibler neighbourhood of radius C, then an analogous definition of admissibility can be given.

First we define the pairwise difference in expected losses of any two actions $(a, a') \in A$ within Γ using:

$$L_{a,a'}(\theta) = L_a(\theta) - L_{a'}(\theta)$$

and the corresponding least favourable pairwise distribution:

$$\begin{aligned}\pi_{\sup}^{(a,a')} &= \arg\sup_{\pi\in\Gamma}\left\{E_\pi[L_{a,a'}(\theta)] : KL(\pi_I \parallel \pi) \leq C\right\} \\ &= Z^{-1}\pi_I(\theta)\exp\left(\lambda[L_a(\theta) - L_{a'}(\theta)]\right)\end{aligned}$$

with expected loss $\psi^{(a,a')} = \int_\theta \pi_{\sup}^{(a,a')}(\theta)[L_a(\theta) - L_{a'}(\theta)]d\theta$.

Definition 1. *We say that an action a is C-dominated, or locally-inadmissible up to level C when,*

$$\sup_C\{C : \psi_{\sup}^{(a,a')} < 0, \quad \forall a' \in A \setminus a\}$$

If a is ∞-dominated then it is globally inadmissible (this retrieves the classical notion of admissibility).

We note that plotting $\psi_{sup}^{(a)}(C), a \in \mathcal{A}$ for values $C \in [0, C_{max}]$ does not give any information as to the admissibility of the actions $a \in \mathcal{A}$. In order to graphically represent admissibility (inadmissibility), it is necessary to look at least favourable distributions defined over the pairwise difference in loss for two actions a, a'. By plotting $\pi^{(a,a')}$ as a function of constraint radius C we can look for actions that are dominated, such that there is no $\pi \in \Gamma$ for which they are optimal.

Predictive tempering as local-minimax solution: Consider the task, or action, to provide a predictive distribution, $\widehat{\pi(y|x)}$, for a future observation y given covariates x. The local proper scoring rule in this case is known to be the self-information logarithmic loss $L(y) = -\log\pi(y|x)$ (Bernardo & Smith, 1994). The conventional Bayesian solution is to report your honest marginal beliefs as $\widehat{\pi(y|x)} = \pi_I(y|x)$, where given a model parametrised by θ we have $\pi_I(y|x) = \int \pi(y|x,\theta)\pi_I(\theta)d\theta$. However this assumes that the model is true and moreover that the prediction problem is stable in time, in that the prediction probability contours do not drift. Both of these assumptions may be dubious, for the latter see the notion of "concept drift" in data mining, e.g. Section 3.1 in Hand (2006). The robust local-minimax solution protects against misspecification and drift and leads to

$$\widehat{\pi^*_{\sup}(y|x)} \propto \pi_I(y|x)^{1-\lambda},$$

for $\lambda \in [0, 1]$. This has the form of tempering the predictive distribution, taking into account additional external levels of uncertainty outside of the modelling framework. In this way, predictive annealing can be seen as a local-minimax action.

Gibbs Posteriors and PAC-Bayesian: Suppose you hold prior beliefs about a set of parameters θ but don't know how to specify the likelihood $\pi(x|\theta)$, and hence lack a model $\pi(x,\theta)$. For example, suppose θ refers to the median of F_X with unknown distribution. Suppose the task (action) is to provide your best subjective beliefs $\pi(\theta|\cdot)$ conditional on information in the data and prior knowledge. We don't have a likelihood but we could have a well defined prior hence $\pi_I(\theta) = \pi(\theta)$. In this situation there may be a well defined loss function on the data that we would wish to *maximise* utility against for specifying beliefs, e.g. for the median we should take

$$L(\theta) = \sum_i |x_i - \theta|$$

The distribution that *minimises* the expected loss within a certain KL divergence of the prior is given by the local-maximin distribution,

$$\pi_a^{(\inf)} = Z^{-1}e^{-\lambda\sum_i |x_i-\theta|}\pi(\theta)$$

This has the form of a Gibbs Posterior or an exponentially weighted PAC-Bayesian approach (Zhang (2006a,b); Bissiri *et al.* (2013); Dalalyan & Tsybakov (2008, 2012)). In this way we can interpret Gibbs posteriors as local-maximin solutions in the absence of a known sampling distribution (Bissiri *et al.*, 2013).

4.1.1 Computing π^*_{sup}, C and $(\psi^{(a)}_{\text{inf}}, \psi^{(a)}_{\text{sup}})$

In most scenarios the Γ-constrained minimax distribution $Z^{-1}\pi_I(\theta)\exp[\lambda L_a(\theta)]$ will not have closed form. Moreover $\pi_I(\theta)$ will often only be represented as a Monte Carlo bag of samples $\{\theta_i\}_{i=1}^m \sim_{iid} \pi_I$. In this case the distribution can be approximated by using π_I as an importance sampler (IS) leading to,

$$\begin{aligned} \tilde{\pi}^*_{\text{sup}} &= \frac{1}{\sum_i w_i}\sum_i w_i \delta_{\theta_i}(\theta) \\ w_i &= \exp[\lambda L_a(\theta_i)] \end{aligned}$$

We can then use $\tilde{\pi}^*_{\text{sup}}$ to calculate $(\psi^{(a)}_{\text{inf}}, \psi^{(a)}_{\text{sup}})$. For a small neighbourhood size and hence small λ relative to $L_a(\theta)$ this IS approximation should be reasonable. In general if π_I is thought to be a useful model to the truth then the neighbourhood size should be kept small. However as λ increases the variance of the importance weights will increase exponentially and the approximation error with it. In this situation the problem appears amenable to sequential Monte Carlo samplers (SMCS) taking $\lambda \geq 0$ as the "time index" although we do not explore this further here. This points to the wider issue of how to choose the size of the neighbourhood Γ. This will of course be highly context dependent and we propose using qualitative explorative methods to investigate the sensitivity of the choice of KL radius C. These are illustrated in section 5.1 and further explained in the discussion in section 6.

4.1.2 Comments on the use of $\Gamma = \{\pi : KL(\pi \parallel \pi_I) < C\}$

Exploring only the best or worst possible outcomes within a given KL ball may be seen as restrictive and unrealistic. The minimax solution assumes a reactive zero-sum game leading to a different solution to any particular action chosen, in conflict with the notion of a stable Nature. To counter this argument we can restrict attention to small neighbourhoods (see the discussion in Chapter 1 of Hansen & Sargent (2008)). An alternative is to look at the distribution of expected losses over all models $\pi \in \Gamma$. However it is not clear how one can enumerate this or compute with the space of distributions as defined by $\Gamma = \{\pi : KL(\pi \parallel \pi_I) < C\}$. It turns out that by changing the neighbourhood from $\text{KL}(\pi \parallel \pi_I)$ to $\text{KL}(\pi_I \parallel \pi)$ a numerical solution exists. However it should be noted that this is at the expense of the coherency argument outlined above.

4.2 $\text{KL}(\pi_I \parallel \pi)$

The change of neighbourhood from $\text{KL}(\pi \parallel \pi_I)$ to $\text{KL}(\pi_I \parallel \pi)$ results in a non-analytic solution to the local-minimax and maximin distributions. However we can use numerical methods to compute both the minimax optimisation and, importantly, the distribution of expected loss within the KL ball. We first consider the numerical solution to $\pi^*_{\text{sup}} = \arg\sup_{\pi\in\Gamma} E_\pi[L_a(\theta)]$, with Γ now defined herein as $\Gamma = \{\pi : KL(\pi_I \parallel \pi) \leq C\}$ for $C \geq 0$.

4.2.1 Approximating π^*_{sup} within Γ

Consider a stochastic representation of π_I via

$$\begin{aligned} \tilde{\pi}_I &\equiv \frac{1}{m}\sum_{i=1}^{m}\delta_{\theta_i}(\theta) \\ \theta_i &\sim \pi_I(\theta) \end{aligned} \tag{4}$$

where θ_i are iid draws from π_I and $m \to \infty$. In practice this is often the actual model that statisticians work with, via a "bag of samples" Monte Carlo representation of π_I. We note for non-degenerate functionals $g(\theta)$ of interest $E_{\tilde{\pi}_I}[g(\theta)]$ converges to $E_{\pi_I}[g(\theta)]$, as $m \to \infty$. To make the solution tractable in defining a KL neighbourhood around π_I we will use the neighbourhood around $\tilde{\pi}_I$. Moreover, in considering the KL divergence between π_I and an alternative model $\pi \in \Gamma$ we will work with a stochastic approximation to π represented as mixtures of the atoms $\{\delta_{\theta_1}, \delta_{\theta_2}, \ldots, \delta_{\theta_m}\}$ in (4),

$$\tilde{\pi} = \sum_i w_i \delta_{\theta_i}(\theta) \tag{5}$$

for $0 \leq w_i \leq 1$, $\sum_i w_i = 1$, where the w_i's can be interpreted as importance weights $w_i \propto \pi(\theta_i)/\pi_I(\theta_i)$, with $E_{\tilde{\pi}}[g(\theta)] \to E_{\pi}[g(\theta)]$, as $m \to \infty$.

The KL divergence between π_I and π can then be approximated via the KL divergence of their stochastic representations,

$$KL(\tilde{\pi}_I, \tilde{\pi}) = \frac{1}{m}\sum_{i=1}^{m}\log\frac{1}{w_i}.$$

with KL ball $\tilde{\Gamma}$ defined similarly, $\tilde{\Gamma} = \{\tilde{\pi} : KL(\tilde{\pi}_I, \tilde{\pi}) < C\}$.

From these definitions, we will now look for the probability measure maximisation

$$\tilde{\pi}^*_{\text{sup}} = \arg\sup_{\tilde{\pi}\in\tilde{\Gamma}} \{E_{\tilde{\pi}}[L_a(\theta)] : KL(\tilde{\pi}_I \parallel \tilde{\pi}) \leq C\} \tag{6}$$

Given the atomic structure of $\tilde{\pi}$ the maximisation (6) leads to a convex optimisation in the weights,

$$\tilde{\pi}^*_{\text{sup}} = \sum_i w_i^* \delta_{\theta_i}(\theta) \tag{7}$$

$$w^* = \arg\sup_w \{\sum_i w_i L_a(\theta_i) : -\sum_i \log(w_i) \leq \lambda, \sum_i w_i = 1\} \tag{8}$$

for which standard numerical methods can be applied.

4.2.2 DP averaging neighbourhood loss

From a Bayesian standpoint its more natural to characterise the variation in expected loss arising from models in the neighbourhood of the approximating model, rather than minimax optimisation. In order to quantify uncertainty or take expectations over distributions in the neighbourhood of π_I we require a probability distribution on a set of probability measures centred on π_I. This is classically a problem in Bayesian nonparametrics, see for example Hjort *et al.* (2010). In particular the Dirichlet Process (DP) can be used to construct a distribution on distributions centred on a baseline π_I.

Definition 2. Dirichlet Process: *Given a state space $\mathcal{X}$ we say that a random measure P is a Dirichlet Process on $\mathcal{X}$, $P \sim DP(\alpha, P_0)$, with concentration parameter α and baseline measure P_0 if for every finite measurable partition $\{B_1, \ldots, B_k\}$ of $\mathcal{X}$, the joint distribution of $(P(B_1), \ldots, P(B_k))$ is a k dimensional Dirichlet distribution $Dir(\alpha P_0(B_1), \ldots, \alpha P_0(B_k))$.*

Using this definition we can then sample from distributions in the neighbourhood of π_I according to,

$$\pi^* \sim DP(\alpha, \pi_I)$$

In practice we consider a draw from the DP via the constructive definition,

$$\begin{aligned} \pi^* &= \sum_{i=1}^{m} w_i \delta_{\theta_i}(\theta) \\ \theta_i &\sim \pi_I \\ \underline{w} &\sim \mathrm{Dir}_m(\alpha/m, \ldots, \alpha/m), \\ m &\to \infty \end{aligned} \tag{9}$$

where the θ_i's are iid from π_I and independent of the Dirichlet weights. This highlights one difficulty with the use of the DP. It is problematic to measure the divergence between any two draws, or from a sample to the baseline measure, as the discrete atomic structure means that two draws do not have the same support. For instance the Kullback-Leibler divergence $\mathrm{KL}(\pi_I \parallel \pi^*)$ is not defined. In order to circumvent this we introduce the notion of coupled-DPs.

Definition 3. Coupled Dirichlet Process. *We say that two or more Dirichlet process samples $\{\pi^{(i)}, \pi^{(j)}\}$ are coupled, $\pi^{(i)}, \pi^{(j)} \sim CDP(\alpha_i, \alpha_j, \pi_0)$ if they share a common set of baseline atoms, $\{\delta_{\theta_s}(\theta)\}_{s=1}^m$ drawn from the baseline measure, π_0, but with independent Dirichlet weights,*

$$\begin{aligned} \pi^{(i)} &= \sum_{s=1}^{m} w_s^{(i)} \delta_{\theta_s}(\theta) \\ \pi^{(j)} &= \sum_{s=1}^{m} w_s^{(j)} \delta_{\theta_s}(\theta) \\ \theta_s &\sim \pi_0 \\ \underline{w}^{(i)} &\sim \mathrm{Dir}_m(\alpha_i/m, \ldots, \alpha_i/m), \\ \underline{w}^{(j)} &\sim \mathrm{Dir}_m(\alpha_j/m, \ldots, \alpha_j/m), \\ m &\to \infty \end{aligned}$$

It is clear that asymptotically $\pi^{(i)}$ and $\pi^{(j)}$ are conditionally independent given the baseline measure as, following the definition of the Dirichlet process, conditionally for every finite measurable partition $\{B_1, \ldots, B_k\}$ we have that $\pi^{(i)}$ is a k-dimensional Dirichlet distribution with parameters $(\alpha_i \pi_0(B_1), \ldots, \alpha_i \pi_0(B_k))$ which are independent of $\pi^{(j)}$. However, using a CDP allows us to characterise a divergence measure between the samples. For example, given two realisations we find the KL divergence as,

$$KL(\pi^{(i)} \parallel \pi^{(j)}) = \sum_{s=1}^{m} w_s^{(i)} \log \left(\frac{w_s^{(i)}}{w_s^{(j)}} \right)$$

Now consider the atomic representation of the baseline measure $\tilde{\pi} \sim DP(\infty, \pi_I)$, so that as in

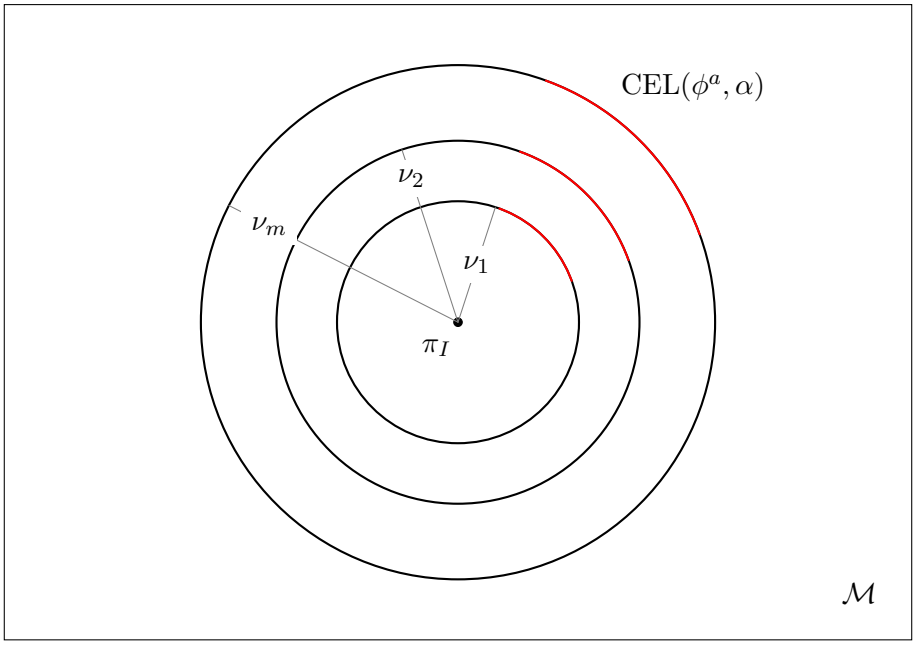

Figure 5: CEL given by DP Sampler with uniform distribution over the KL divergence radii.

(4),

$$\begin{aligned} \pi^* &= \sum_{i=1}^{m} \frac{1}{m} \delta_{\theta_i}(\theta) \\ \theta_i &\sim \pi_I \\ m &\to \infty \end{aligned} \tag{10}$$

Then we can compute a measure of divergence to the baseline measure using a coupled Dirichlet process, $\tilde{\pi}^*, \tilde{\pi}_I \sim CDP(\alpha, \infty, \pi_I)$, for which,

$$\begin{aligned} KL(\tilde{\pi}_I \parallel \tilde{\pi}^*) &= \frac{1}{m} \sum_{i=1}^{m} \log\left(\frac{1}{m w_i}\right) \\ \underline{w} &\sim \text{Dir}_m(\alpha/m, \ldots, \alpha/m), \end{aligned}$$

The KL is itself a random variable and it is interesting to characterise its distribution from the perspective of the neighbourhood size around π_I. To investigate this it will be helpful to consider the representation of the Dirichlet as a normalised sum of Gamma random variables,

$$\begin{aligned} w_i &= \frac{G_i}{G_1 + \cdots + G_m} \\ G_i &\sim \text{Gam}(\alpha, 1) \end{aligned}$$

with $m \to \infty$ and let $G = \sum_i G_i$, and hence $G \sim \text{Gam}(m\alpha, 1)$. Under this representation we have for $\tilde{\pi}_0, \tilde{\pi}^* \sim CDP(\infty, \alpha, \pi_I)$,

$$KL(\tilde{\pi}_0 \parallel \tilde{\pi}^*) = \log(G) - \frac{1}{m} \sum_i \log(G_i) - \log m. \tag{11}$$

Proposition 1. *Let $\tilde{\pi}^*, \tilde{\pi}_I \sim CDP(\alpha, \infty, \pi_I)$ then π^* is distributed on the outer shell of a KL ball, $\tilde{\Gamma}$, centered on $\tilde{\pi}_I$ of radius ν with,*

$$\nu := g(\alpha) := \log \alpha - \psi_0(\alpha)$$

where $\psi_0(\cdot)$ denotes the digamma function.

Proof. From the properties of the log-Gamma random variable we have $X \sim \log[\text{Gam}(a,b)]$ then $E[X] = \psi_0(a) + \log(b)$ and $Var[X] = \psi_1(a)$, where ψ_1 is the trigamma function. Hence in (11) $E[\log(G)] = \psi_0(m\alpha)$, which for large $m \to \infty$ we have $\psi_0(m\alpha) \to \log(m\alpha)$ and $\frac{1}{m}\sum_i \log(G_i) \to \psi_0(\alpha)$, leading to, $E[KL(\pi \parallel \pi^*)] = \log(\alpha) - \psi_0(\alpha)$ and,

$$\begin{aligned} \text{Var}[KL] &= \text{Var}[\frac{1}{m}\sum_i \log(G_i)] - \text{Var}[G] \\ &= \frac{1}{m}\text{Var}[\log G_i] - \text{Var}[G] \\ &= \frac{1}{m}\psi_1(\alpha) - \psi_1(m\alpha) \end{aligned}$$

following from the linear dependence of G and $\sum_i G_i$, from which we see that $\text{Var}[\text{KL}] \to 0$ as $m \to \infty$. The result follows. □

This motivates statistics analogous to those defined in section 3. For instance, we can define a prior $P(\nu)$ over the KL radius $\nu \in [0, C]$ to induce a distribution over measures $\pi \in \Gamma$ and therefore a new distribution $F^{\Gamma}_{Z_a}$ over losses z (note this new distribution of loss will have the same support as the original F_{Z_a}). Sampling from this distribution can be done by sampling $\nu \sim P(\cdot)$, and then $\theta_i \sim \pi_I$, with weights $w_i \sim \text{Dirichlet}(\alpha)$, where $\alpha = g^{-1}(\nu)$. Reiterating this scheme, allows for calculation of statistics such as the Conditional Expected Loss (CEL) value, but now with respect to the prior set of distributions $\pi \in \Gamma$. A pictorial representation is given in figure 5. The Γ-conditional expected value of a decision is defined as an average of the expected values of the $q\%$ least favourable distributions in Γ at exactly a KL divergence of ν from π_I. These are drawn in red.

5 Applications

5.1 Optimal Screening Design for Breast Cancer

Public health policy is an area where the application of statistical modelling can be used to optimally allocate resources. Breast cancer screening for healthy women over a certain age to detect asymptomatic tumours, is a hotly debated and controversial issue for which it is difficult to fully quantify the benefits. A recent independent review (Marmot *et al.* (2012)), commissioned by Cancer Research UK and the Department of Health (England) concludes that only a randomised clinical trial would fully resolve this issue. This is of course the gold standard which permits causal inference. However a primary issue is determining the optimal screening schedule, consisting of a starting time t_0 (age of first screen), and a frequency δ for subsequent screens. It is of course sharply infeasible to trial all combinations of schedules (t_0, δ). An optimal trial design however can be constructed using a statistical model of disease progression throughout a population. Parmigiani (1993) proposed using a semi-Markov process consisting of four states which generalises to any chronic disease characterised by an asymptomatic stage. All individuals start in state A, disease-free. They then transition either to the absorbing state D (death) or contract the disease, modelled by a transition to state B, the pre-clinical stage. This is followed by a transition to either the clinical stage of the disease or death. It was assumed that each transition happens after a time t with the following densities:

$$t_D \sim h(t|c,d) = \text{Weibull}(\alpha, \beta)$$

$$t_B \sim f(t|\mu,\sigma^2) = \text{LogNormal}(\mu, \sigma^2)$$

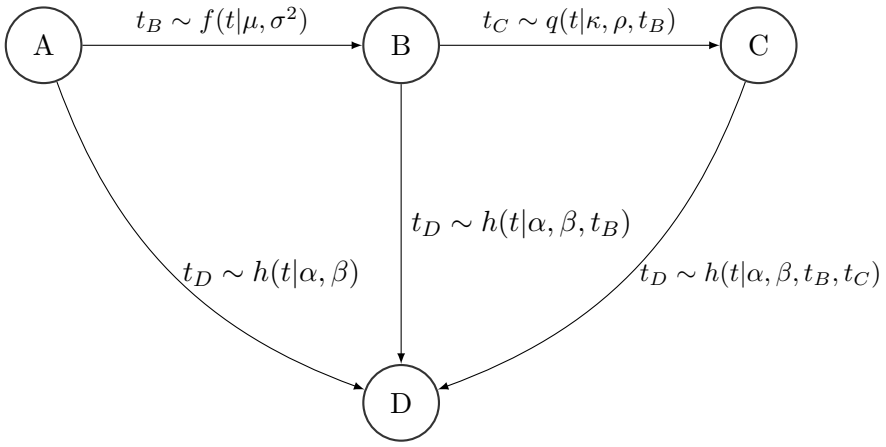

Figure 6: Graphical model of the transitions and transition densities between states.

$$t_C \sim q(t|\kappa, \rho) = \text{LogLogistic}(\kappa, \rho)$$

Figure 6 shows a graphical model of the four state semi-Markov process with transition densities. An individual is characterised by the triple $t = (t_B, t_C, t_D)$ where the symptomatic stage of the disease is contracted only when $t_D < t_B + t_C$ (assuming that all individuals will contract the disease if they lived long enough). For a screening schedule $a = (t_0, \delta)$ the loss function is defined as follows (a function of the times $t = (t_B, t_C, t_D)$):

$$L(a, t) = r \cdot n_a(t) + \mathbb{1}_C \tag{12}$$

where n_a is the number of screening schedules an individual will receive during their lifetime, until they die or enter into the asymptomatic stage of the disease. $\mathbb{1}_C$ is the indicator function, taking value 1 for the event that the pre-clinical tumour is not detected by screening or occurs before t_0, and zero otherwise. r trades off the cost of one screen against the cost incurred by the onset of the clinical disease. Each screen has an age-dependent false-negative rate modelled with a logistic function:

$$\beta(t) = \frac{1}{1 + e^{-b_o - b_1(t - \tilde{t})}}$$

where $\tilde{t}$ is the average age at entry in the study group. To simulate transition times for individuals from this model, we used 2000 posterior parameter samples for $\theta = (\mu, \sigma^2, \kappa, \rho, b_0, b_1)$ given in the supplementary materials of Wu *et al.* (2007). This is based on data from the HIP study Shapiro *et al.* (1988). Figure 7 gives an example of 10^4 sampled times along with a histogram of losses corresponding to the schedule $a = (t_0 = 40, \delta = 0.75)$ (units in years) and a trade-off parameter $r = 10^{-3}$. In the model we calibrate the Weibull distribution with values $\alpha = 7.233, \beta = 82.651$ (the values used by Parmigiani (1993)).

The losses incurred for a particular schedule $a = (t_0, \delta)$ can be seen to be highly bimodal. Most of the population do not contract the disease and therefore contribute a loss of $r \cdot n_a$ (cost of screen times number of total screens during lifetime). The loss contributed by those who do contract the clinical stage of the disease is of magnitude $\frac{1}{r}$ greater by definition. We considered 12 alternative schedules, consisting of all combinations of starting ages taken from the set $(40, 45, 50)$ (in years) and screening frequencies of $(9, 12, 15, 18)$ (in months). This choice of screening schedules is somewhat illustrative for our purposes: the optimal schedule will heavily depend on the choice of r (trade-off ratio in (12)) which we do not attempt to estimate (the value

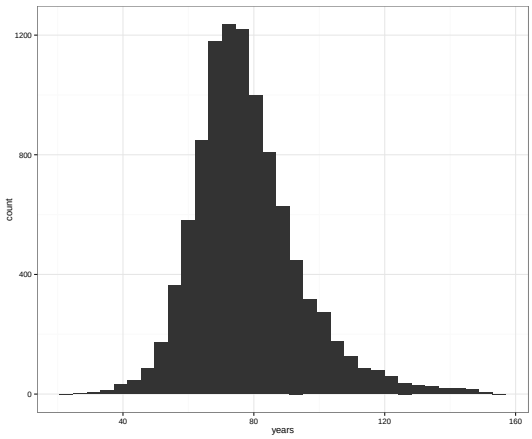

(a) Transition times from A to B generated from model (in years).

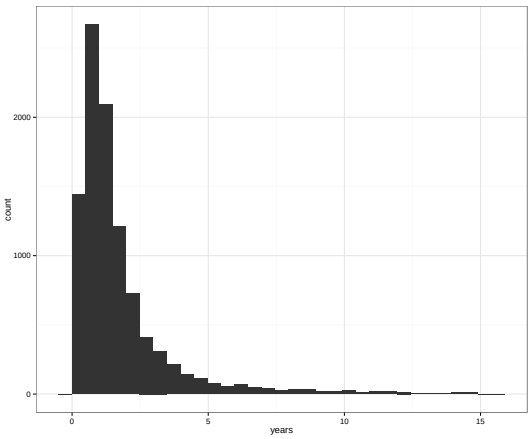

(b) Sojourn times: transition times from B to C.

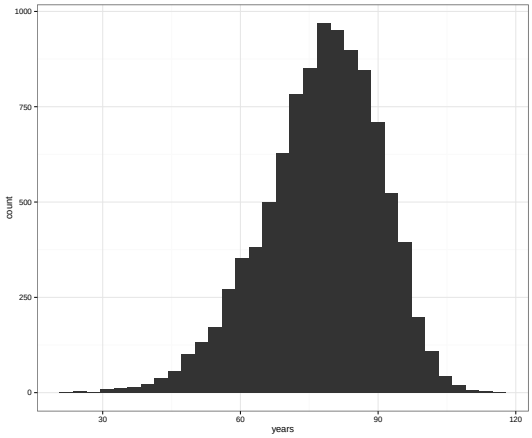

(c) Lifetimes generated from Weibull$(7.2, 82.6)$.

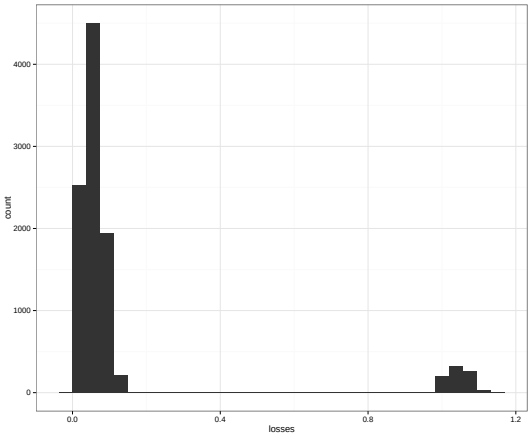

(d) Histogram of losses with schedule $(40, .75)$.

Figure 7: Summary of 10^4 transition times generated using the posterior parameter samples from Wu *et al.* (2007) with losses reported under a screening schedule with starting age of 40 years and frequency of nine months (0.75 years).

10^{-3} was taken from the section 4.5 of Ruggeri *et al.* (2005) where the authors also consider this application).

Figure 9 illustrates the diagnostic plots for the VaR, CVaR and CER statistics defined in section 3 with the schedules (decisions) aforementioned. Figure 9d clearly indicates the heavy downside of every breast cancer screening schedule: almost all of the expected loss is being contributed by events coming from the first 10^{th} to the first 20^{th} quantile of loss, depending on the schedule. It also demonstrates the influence of the choice of age of first screening versus the choice of screening frequency. However, when considering the relationship between the loss function and the density estimates of the samples, we see that although the high loss events are infrequent, their density estimates are uniformly distributed relative to all events. This is as discussed in section 3.4, where we argue that the relationship between the loss function and the density estimate of the samples can help decide whether the expected loss estimate is being driven by a few events with very low posterior measure, or whether the high loss events are evenly spread out over the posterior predictive. In figure 8, we plot the relationship between the loss function and the density estimate of the corresponding transition times (estimated using the function *npudens* from the R package *np*, with default settings). The first three sub-figures show this for each marginal density (transition times t_B, t_C and t_D) and the fourth sub-figure plots this

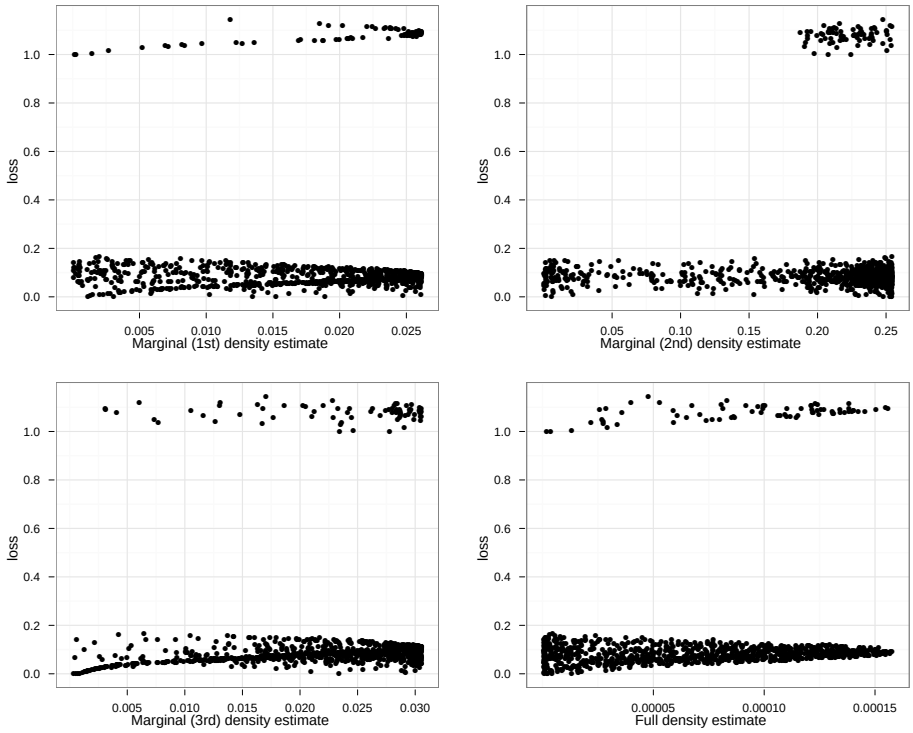

Figure 8: Loss-density plot as explained in section 3.4. The first three sub-figures show the loss-density relationship for each marginal distribution (over times t_B, t_C, t_D) and the last plot shows the relationship for the full multivariate distribution over transition times.

for the full multivariate distribution. From these plots we can conclude that there is stability in the approximation of the expected loss estimate, with no dependence between the density estimate and the loss.

In order to explore how the optimality of different schedules changes with varying degrees of misspecification (i.e. radius of Kullback-Leibler ball around the model π_I) we plot the difference in expected loss under the local-minimax for all schedules against the optimal schedule under π_I (start = 40, frequency = .75). Figure 10a plots this difference for the 12 schedules. There are multiple crossing points, with the final schedule (50 years, 12 months) ending as most optimal for KL distance greater than ≈ 5.2. Figure 10a would suggest that the model is not decision-robust with respect to model misspecification. However, one primary question is whether or not a KL distance $\nu = 5.5$ is plausible or not in this context. One method for calibrating 'feasible' values of ν is to visually inspect the univariate marginals of π^*_{sup}. The loss function is defined over the transition times $\{t_B, t_C, t_D\}$ (as shown in figure 6) so we plot in figures 10b and 10c the marginal density estimates of the least favourable distributions of schedule 1 (optimal under π_I) and schedule 10 (optimal under constrained minimax for $\nu \geq 5.5$), and of schedules 1 and 2 (optimal under constrained minimax for $\nu \geq 1$) respectively, overlaid on marginal densities of π_I. These marginal densities can be directly assess as to their plausibility. It would seem that for this particular application, decision optimality is highly sensitive to the parametrization of the model. These methods also allow for a better understanding of the effect each transition time has on the loss function (which is a complex function of all three). If such models are used to help inform health policies then such considerations may be extremely important.

It is also possible to directly plot the values of the weights $\{w^*_i\}_{i=1}^n$ that define the approximation to the local-minimax distribution for the schedules 1&2, and 1&4 at KL divergences of

1 and 5 respectively. This graph answers two questions. Firstly, is sample size an issue when approximating the least favourable distribution for a given KL divergence? Secondly, where is the model most sensitive to misspecification? If the sample size is too small, then as the KL radius is increased, all the weight will be put onto a few samples with highest loss. However, for values of C, the KL radius, for which this is not the case, these weights can be projected onto the marginals (that are included in the loss function) and inform the user where the model is most sensitive to misspecification. Figure 11 plots this for both pairs of schedules and KL divergences highlighted by the optimality analysis represented in figure 10a. We observe that the most concentrated alteration happens to the transition times t_C, with the least favourable model favouring shorter periods of pre-clinical cancer. This is of course to be expected, as this directly effects the possibility of detection by scanning.

As well as plotting the expected losses under the local-minimax distributions given by the analytic results in section 4.1, we can also exploit the theory developed in section 4.2.2 concerning Dirichlet Processes centred at π_I with concentration parameter α. Repeatedly drawing weights from a symmetric Dirichlet distribution with varying concentration parameter α, we estimate the local-minimax loss for different values of KL radius C (using the approximation $C = \log(\alpha) - \psi(\alpha)$, where ψ is the digamma function). This is shown in figure 12a. We note that the algorithms used to generate the Dirichlet weights (stick breaking process, generating beta weights) run into difficulties for KL values less than 0.05 ($\alpha \geq 17.5$). That is to say a tolerable level of convergence was not reached ($\epsilon = 10^{-4}$) given the number of samples available. Stable estimates of the local-minimax loss necessitated a high number of reruns (order of 10^4). However for the Γ-CEL(q), only (1-q)% of the samples are discarded which leads to faster convergence to a stable estimate. This is shown in figure 12b, for $q = .1$. Further work is needed to make these methods more user friendly.

As a final sensitivity analysis, we use the notion of *C-domination*, see definition 1. Recall this takes a composite loss function, defined as the difference between the losses of two actions a_1 and a_2 for a given $\pi \in \Gamma$. Plotting the value of the local-minimax loss for this composite loss as a function of the KL radius will determine balls Γ_C for which there does not exist a $\pi \in \Gamma_C$ such that a_2 is preferable to a_1. That is to say the expected loss is always negative within Γ_C. Figure 13 plots this for a selection of 4 schedules (starting age of 40 and frequencies of 9, 12, 15 and 18 months). We note that the first sub-figure shows the schedule (40,9) dominates the three others for KL-radii of $\approx .3 * 10^{-3}$, $2 * 10^{-3}$, and $3.2 * 10^{-3}$ respectively. This reinforces the fact that this decision problem is not stable, as the optimal schedule only dominates the other alternatives for very small balls Γ_C (where 'small' is defined relatively here, using diagnostic plots such as figure 10).

This application highlights an interesting distinction that must be made when considering model misspecification in a decision-theoretic context. The loss surface is very flat for changes in screening schedule. That is to say, there is little relative difference in expected loss between similar screening schedules. This is also noted by Ruggeri *et al.* (2005) in their analysis of the problem. Interestingly, this property also seems to hold as we expand the KL ball around the model π_I (see figure 10a), although the optimal schedule changes. This particular application is robust to changes in the model (in an expected loss sense) but not however decision robust. I.e. small perturbations to the model will change the optimality of an action a^*. We discuss this idea further in section 6.

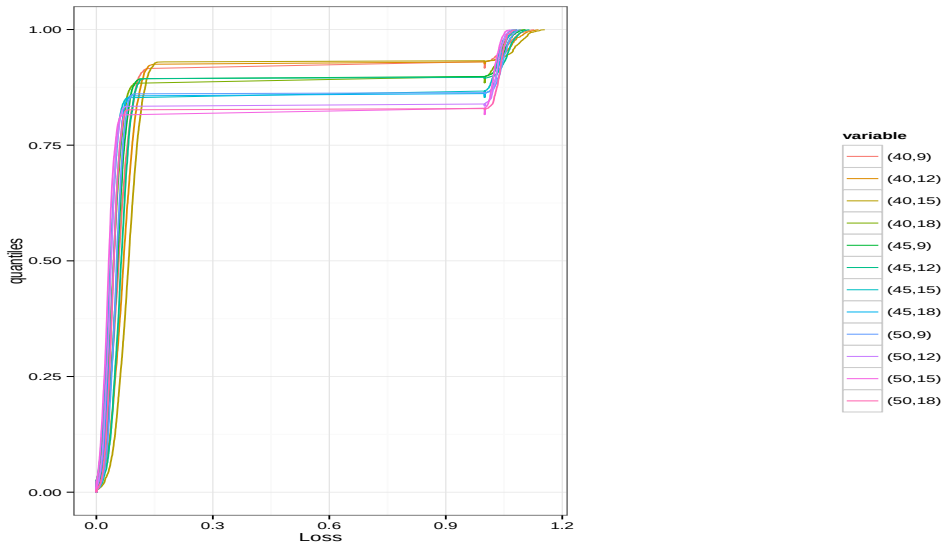

(a) Loss distribution with legend.

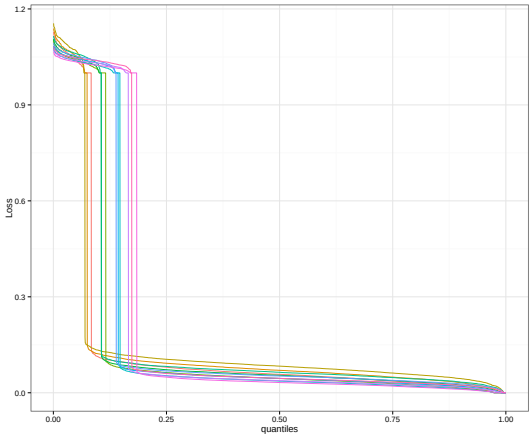

(b) Sorted losses: inverse distribution of loss.

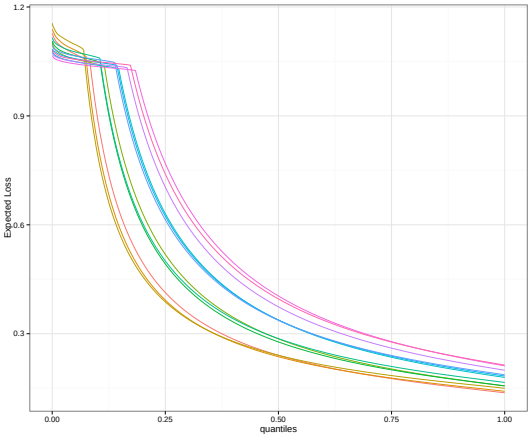

(c) Conditional Value at Risk.

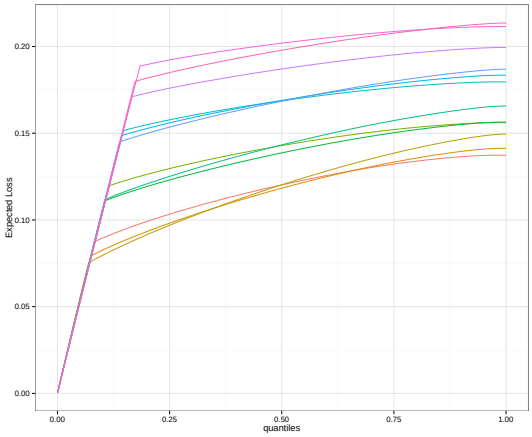

(d) Cumulative Expected Return.

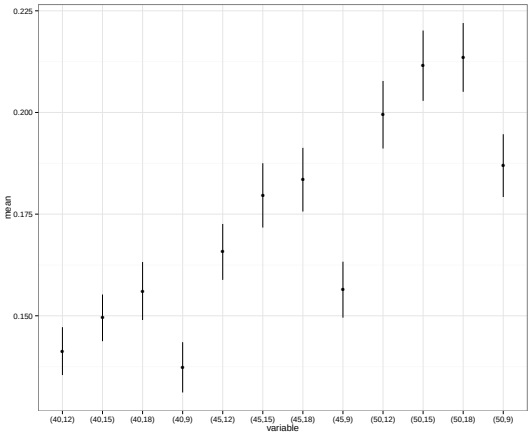

(e) Interval of two standard deviations centred at expected loss.

Figure 9: Summary of losses generated using the posterior samples from Wu *et al.* (2007) and 12 different schedules.

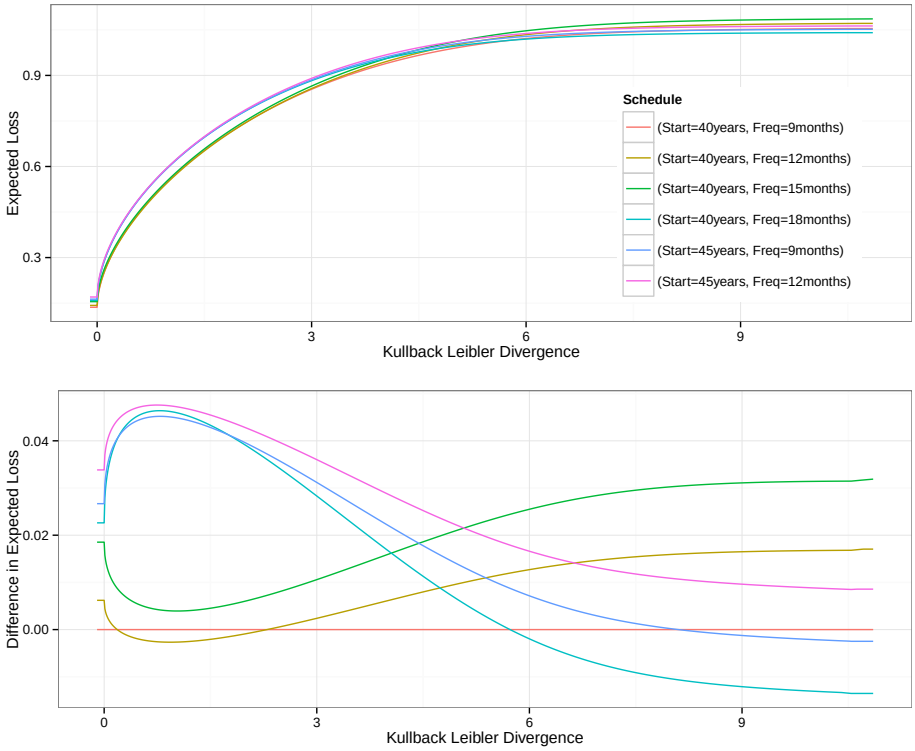

(a) The top figure plots $\phi^{(a)}_{sup}(\nu)$ for six different actions (given in legend). The bottom figure plots the difference of these values with respect to the most optimal under π_I(1: 40 years, 9 months) and therefore shows the change in optimality as the KL radius increases. Note for a KL radius greater than ≈ 5.5 schedule 1 is suboptimal.

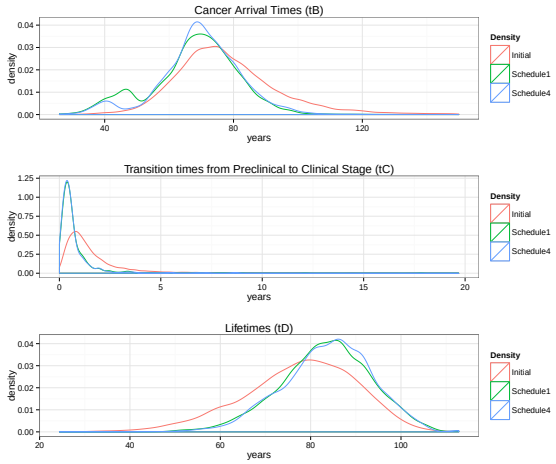

(b) Least favourable densities for radius $\nu = 5$

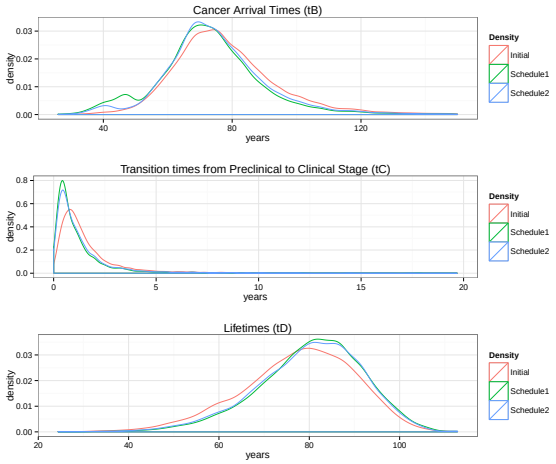

(c) Least favourable densities for radius $\nu = 1$

Figure 10: Interpreting Kullback-Leibler divergence: marginal densities over Θ. Figures (b) and (c) show the density estimates of the three marginal transition times (as explained in figure 6). In red is the density estimate given by the model π_I, and we overlap density estimates of the least favourable distributions for schedules 1 (green; start = 40 years, frequency = 9 months) and 10 (blue; start = 50 years, frequency = 12 months) at a Kullback-Leibler divergence of 5.5 (approximately the crossing point in figure 10a where schedule 1 becomes suboptimal) for figure (b). Same for figure (c) but with schedules 1 and 2 (blue; start = 40 years, frequency = 12 months)

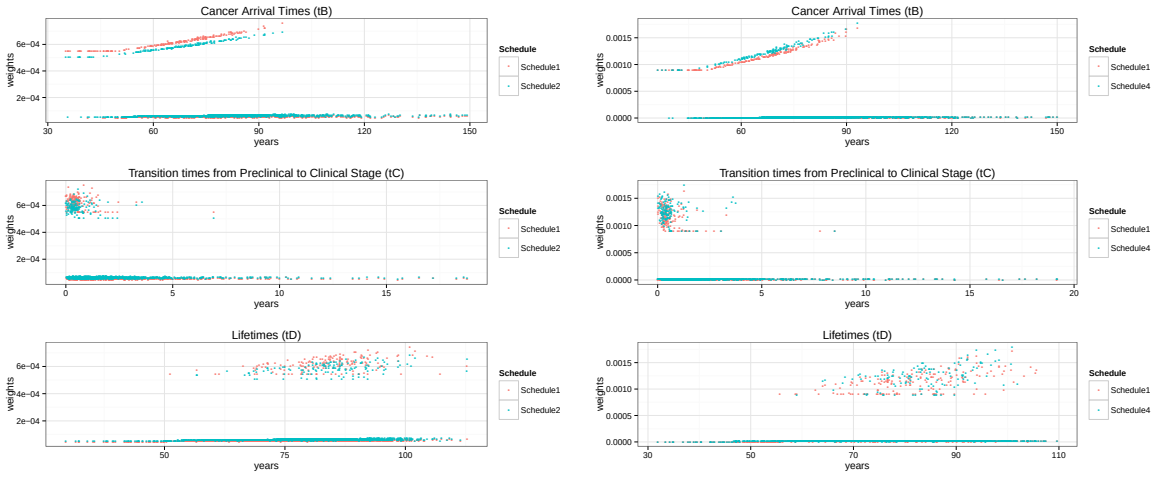

(a) (b)

Figure 11: Plotting the weights corresponding to the marginals shown in Figure 10. This allows a better understanding of the dependence between the least favourable distributions and the model of transition times. Figure (a) corresponds to the local-minimax weight plot for schedules 1 and 2 at KL divergence $C = 1$, and figure (b) corresponds to the local-minimax weight plot for schedules 1 and 4 at KL divergence $C = 5$.

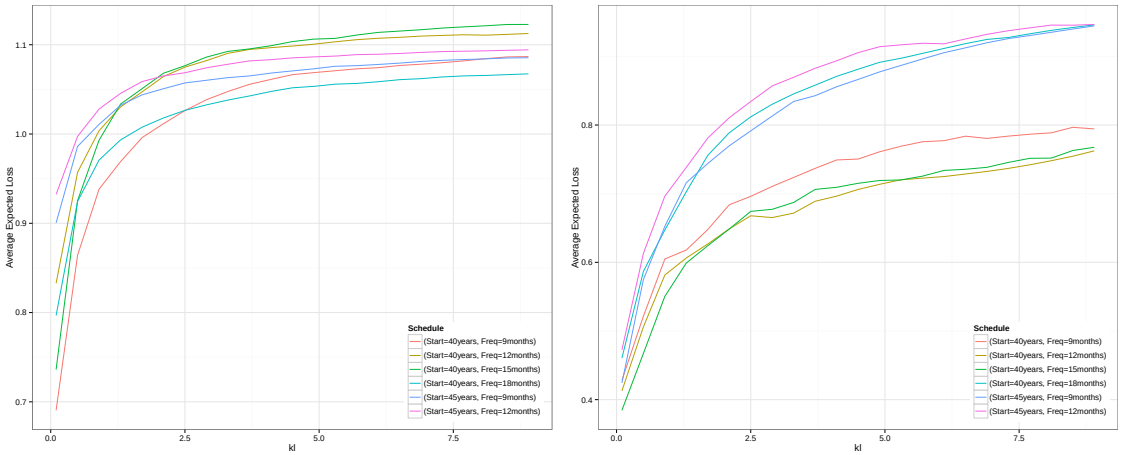

(a) DP sampler: Least favourable distribution (b) DP sampler: 10% CEL as a function of KL divergence.

Figure 12: Possible statistics given by DP(π_I, α) sampler, with $x-$axis shown in KL divergence $(log(\alpha) - \psi(\alpha))$

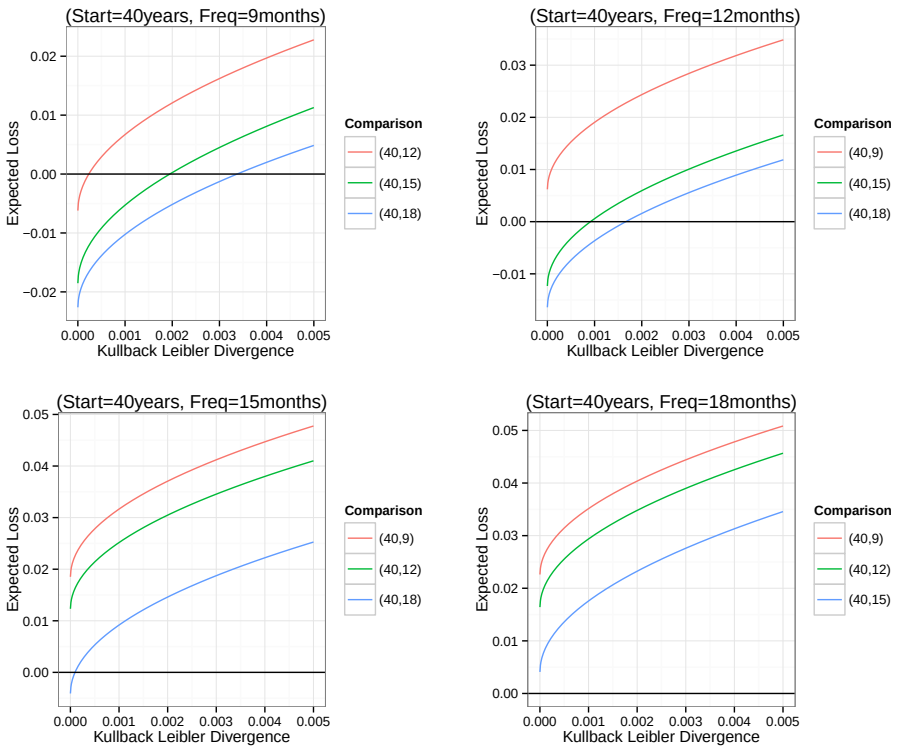

Figure 13: Least favourable distributions when comparing two actions. Each plot compares one schedule with the three others taken from set for starting age of 40 and frequencies of 9, 12, 15 or 18 months. A crossing point with the y-axis indicates the minimum Kullback Leibler divergence necessary for the action to be less favourable than its counterpart.

5.2 Bayesian Variable Selection with Cost

For high-dimensional parameter spaces, one goal of statistical inference is to select an optimal subspace that reduces dimensionality but permits informative modelling. Regression is a standard example where the number of potential predictors p for a response y can be large. We consider the case when data collection comes at a cost, i.e. it is an increasing function of the number of covariates included. Therefore it is necessary to find a minimal set of maximally informative variables. This problem is only well-posed if approached in a decision-theoretic context. If $\mathcal{X} = \{x_j\}_{j=1}^{p}$ is the set of possible regressors, then an action a is defined as a subset $\gamma_a = \{\gamma_i\}_{j=1}^{p}$ where $\gamma_j = 1$ if the j^{th} regressor is included and zero otherwise, see Brown *et al.* (1999). γ_a is a possible model, and the decision task is one of model selection. The loss function defined over a model γ_a must trade accuracy of prediction against total cost of data collection:

$$L(\gamma^a) = L_p(\gamma^a) + C \sum_{j=1}^{p} c_j * \gamma_j^a \tag{13}$$

Here, L_p is the loss incurred when predicting $\pi(y|X, \gamma_a)$, and c_j is the cost of including the j^{th} variable in the regression. We make the assumption that the cost is constant with respect to the inclusion of each variable. C is a constant that represents the conversion of units of cost into units of expected predictive loss. The calibration of this parameter is an immediate issue, but this is context dependent and necessitates expert judgement as does the correct use of any statistical model.

To illustrate model robustness under misspecification in this particular context, we look at data collected from a large study by the RAND corporation on quality of hospital care (QoHC) in the US (1980s), Kahn *et al.* (1990); Draper *et al.* (1990). It is typically very expensive to directly assess QoHC. However in theory it should be possible to approximate it by considering each hospital to be a black box and measuring the inputs and outputs. For example, comparing mortality rates of different hospitals (outputs) based on sickness at admission (inputs). The problem is therefore reduced to selecting the most reliable measure of sickness at admission (disease specific) whilst also taking into account their relative costs of measurement.

The RAND dataset consists of 83 regressors and a univariate response $y \in \{0, 1\}$ which is 1 if the patient dies within 30 days of admission and zero otherwise. The sample is of size $n = 2,532$, all elderly American patients hospitalised with pneumonia in the period between 1980-1986. Because of high redundancy in many of the predictors, the aim of the study was to select a minimal set of variables that would best predict the outcome y. Standard classical backward-selection methods were initially used to choose a parsimonious set of $p = 14$ variables with good predictive accuracy. However, as noted in Fouskakis & Draper (2008), this does not take into account the difference in the cost of collection for the 83 variables. This cost is estimated in minutes (time needed on average to extract the corresponding measure from medical records) and varies from 0.5 to 10 with median cost 1.

Fouskakis & Draper (2008); Fouskakis *et al.* (2009) look at principled methods for trading off cost of variable inclusion against loss in predictive accuracy. The problem is posed as logistic regression, and we choose the logarithmic scoring rule to compute the expected loss under a model γ_a, defined as:

$$L(r, i) = -\ln r_i$$

where r is a the vector of probabilities reported, when the i^{th} event happens. The true predictive

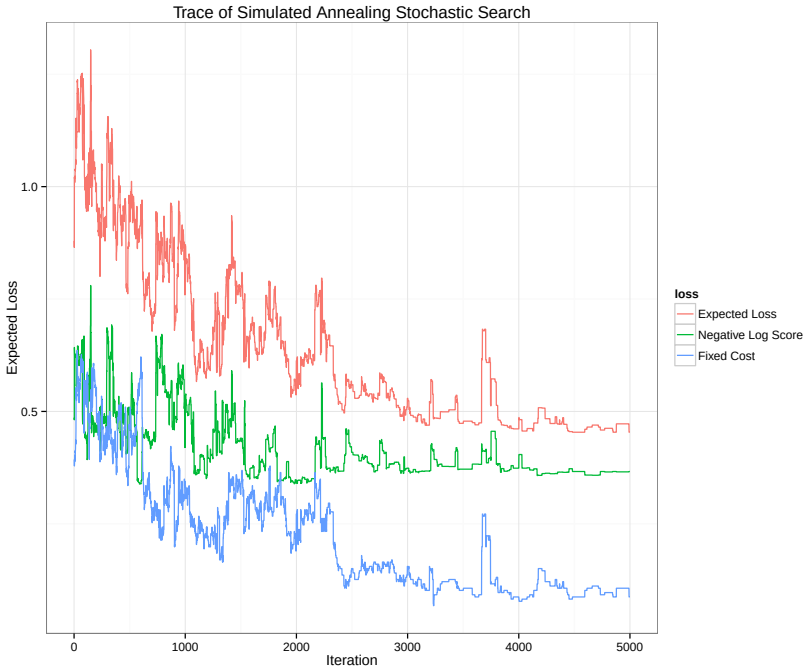

Figure 14: Trace of Expected Loss during a typical run of our simulated annealing search algorithm

distribution over outcomes given regressors x_i is unknown, but we can use the full model $\pi_I(y|X)$ (built on all available covariates) as our best approximation to Nature. This is the marginal of the logistic regression parameters, i.e.:

$$\pi_I(y|X) = \int_\beta \pi(y|\beta, X)\pi(\beta)\mathrm{d}\beta$$

where the integral is approximated using Monte Carlo sampling. This was done using the R package BayesLogit developed by Polson *et al.* (2013)[4]. Therefore the expected loss from diminished predictive accuracy by using the reduced model $\pi_{\gamma_a}(y|X)$:

$$E_{\pi_I}[L_p(\gamma_a)] = (-1) * [\pi_I(y = 1|X) \ln \pi_{\gamma_a}(y = 1|X) + \pi_I(y = 0|X) \ln \pi_{\gamma_a}(y = 0|X)]$$

The space of possible models γ is of size $2^{83} \approx 9.7 \times 10^{24}$, which with modern computational resources is too large for brute force methods. In order to select a set of 'good' models, we implement a simulated annealing search over the space $\{0, 1\}^{83}$, starting with 10 random positions initialised at 1[5]. New proposals are generated by flipping the value of a random position with the acceptance rate for higher expected loss models following a geometric cooling regime. Fouskakis & Draper (2008) have very thoroughly explored stochastic optimization methods for variable selection on this dataset. Our main interest is to find a set of plausible solutions on which to test robustness properties under assumptions of model misspecification. We chose the top 6 actions from 10 runs of the algorithm (5000 iterations, random starting vector with 10 non zero entries). These are fully listed in table 1, appendix B, along with the corresponding costs and names. A complete table of all the available variables can be found in Fouskakis &

[4]We used the default settings of the method *logit*, but drawing 2000 samples (samp=2000) and burn-in of 10000 samples (burn=10000).

[5]We choose to initialise 10 variables for the starting vector, as the results from Fouskakis & Draper (2008) indicate a peak in expected utility as a function of the number of predictors retained between 6-10 variables.

Draper (2008). Figure 14 shows the trace of the expected loss of actions searched during a typical run using this implementation on the RAND dataset. The expected loss value (shown in red) is broken down into its two components, the contribution of the negative logarithmic score (shown in green) and the contribution from the fixed variable cost (shown in blue). The cost vector (units given in minutes) was normalised so that $\sum_{j=1}^{83} c_j = 1$. The constant C in equation (13) was chosen to be 1.

Given a choice of model γ_a, under a negative logarithmic scoring rule, the minimax value of the decision is

$$\max_{y'=\{0,1\}} \{-\ln \pi_{\gamma_a}(y'|X)\} = -\ln\left(\min_{y'=\{0,1\}} \{\pi_{\gamma_a}(y'|X)\}\right)$$

That is to say, the logarithmic score when the more unlikely event occurs. An intermediate least favourable distribution π^*_λ, using the notation defined earlier, for $\lambda \geq 0$ is:

$$\pi^*_\lambda(y'|X) \propto \pi_I(y'|X) e^{\lambda L_p(\gamma_a, y')}$$
$$= \pi_I(y'|X) e^{-\lambda \ln(\pi_{\gamma_a}(y'|X))} = \pi_I(y'|X)\pi_{\gamma_a}(y'|X)^{-\lambda}$$

We cannot evaluate this over future outcomes, but this can be approximated by bootstrapping X values from the empirical $\hat{F_X}$ and y's from the predictive $\pi_I(y|X)$. I.e. for $i = 1, .., n$, computing the value of $\pi^*_\lambda(y_i|X_i)$ for a specific value of λ. The approximate Kullback-Leibler divergence of this distribution π^*_λ is then given by:

$$KL(\pi^*_\lambda || \pi_I) = \frac{1}{n}\sum_{i=1}^{n}\left(\sum_{y'=1,2} \frac{\pi_I(y'|X_i)\pi_{\gamma_a}(y'|X_i)}{Z} \log \frac{\pi_{\gamma_a}(y'|X_i)^{-\lambda}}{Z}\right)$$

where Z is the normalization constant for the distribution π^*_λ at (y_i, X_i). Figure 15 shows the minimax expected loss of 6 actions chosen by the stochastic search algorithm as a function of the Kullback Leibler radius of ball surrounding $\pi_I(y|X)$. Values for the KL radius less than zero represent the expected loss under the $\mathcal{D}_{closed}$ model (this allows for easier comparison of actions in the plot). It is interesting to note the complete reversal of optimality as the size of the relative entropy ball increases. The ranking of actions under minimax expected loss (increasing) for balls greater than approximately 0.025 is the reverse ranking under the $\mathcal{D}_{closed}$ model, and corresponds to the ranking given by the entropy of the predictive distributions π_{γ_a}. That is to say, highest entropy corresponds to smallest minimax expected loss for balls Γ_ν with ν larger than approximately 0.025. The negative entropy of the predictive distribution provides a measure of confidence. A distribution with entropy near zero incurs higher losses under the minimax action as it predicts an outcome with greater certainty, the extreme case being the predictive distribution with probabilities $(0, 1)$. Over-confidence in predictions can be a consequence of misspecification and is penalised by the minimax expected loss.

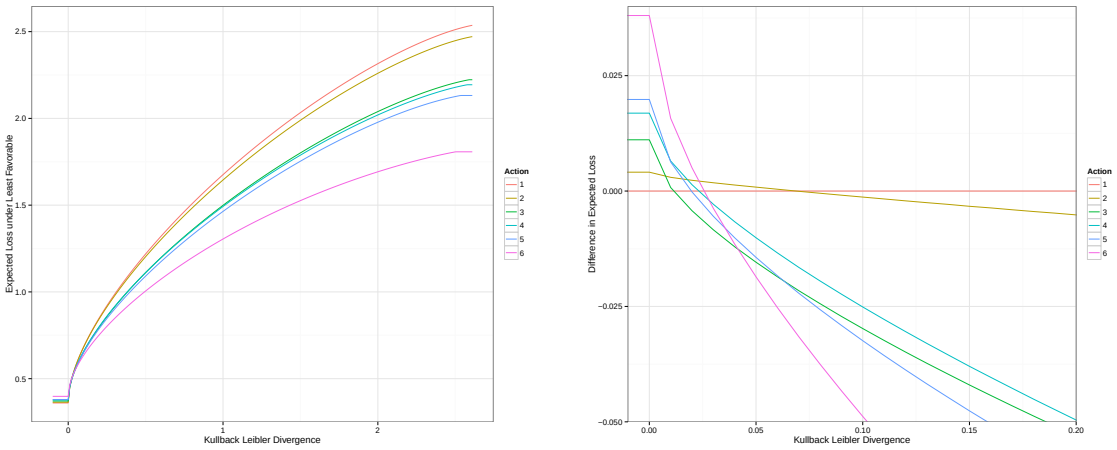

(a) Expected Loss under least favourable distribution.

(b) Zoomed in at origin: showing crossing point.

Figure 15: Comparison of optimality under minimax expected loss as a function of KL radius, for the top 6 actions selected by simulated annealing search. We plot the difference in minimax expected loss in reference to the optimal action in the $\mathcal{D}_{closed}$ decision-system (action 4). Note that this action becomes the least optimal for KL radii greater than 0.1 (crossing points approximately between 0.0125 and 0.075).

6 Conclusions

The goal of this article is to aid decision makers by providing statistical methods for exploring sensitivity to model misspecification. We hope this will generate further debate and research in this field. The increase in complex high-dimensional data analysis problems, "big-data", has driven a corresponding rise in approximate probabilistic techniques. This merits a reappraisal of existing diagnostics and formal methods for characterising the stability of inference to known approximations.

In one sense, the formal methods presented in Section 4 can be considered as extensions of the conditional Γ-minimax approach (Vidakovic (2000)). However, we advocate that the neighbourhood should be defined with respect to the marginal distribution on only those elements that enter into the loss function. Moreover we have shown that the Kullback-Leibler divergence is the only coherent measure to use for Γ in local-minimax updating. Further motivation for using KL is given in Chapters 1 and 9 in Hansen & Sargent (2008). To this we might add, that $\mathrm{KL}(p \parallel q)$ is invariant to re-parametrisation; that in information theory it represents the number of bits of information needed to recover p from model q; it represents the expected log-loss in using q to approximate p when using proper local scoring rules (Bernardo & Smith, 1994); and KL bounds the L1 divergence $\mathrm{KL}(p \parallel q) \leq \parallel p - q \parallel_1$. However, none of this provides a constructive approach for choosing the KL radius C. In chapter 9 of Hansen & Sargent (2008), the authors suggest using detection error probabilities to calibrate the size of the neighbourhood Γ. This stems from the concept of statistically indistinguishable models given a finite data sample of size N. Using model selection principles based on likelihood ratio tests, the user determines a plausible probability (a function of the radius C) of selecting the wrong model given the available data, and then inverts this value to find C (by simulation). Although this is a principled method, in many cases even the detection error probability could be difficult to calibrate. Here we propose a more explorative approach by investigating the marginal densities to visualise the differences between π_I and π_{sup}, as these densities are often highly interpretable (for an illustration see

figure 10 in section 5.1). This is an area that warrants further consideration.

In terms of implementation, we showed that the formal approaches have simple numerical solutions via re-weighted Monte Carlo samples drawn from the approximating model; using exponentially tilted weights for the local-minimax solution and stochastic Dirichlet weights for the marginal loss distribution over Γ.

To complete this discussion of robustness under model misspecification, it is important to note the distinction between "decision robustness" and "loss robustness" as discussed in Kadane & Srinivasan (1994). A system is said to be decision robust if perturbations to the model do not effect the optimality of an action $\hat{a}$. On the other hand, it is said to be loss robust, if those perturbations do not effect the overall expected loss of the action $\hat{a}$ (in a relative sense). It is clear that a decision system can have one property without the other. Which is more desirable will be highly context dependent. Throughout the article we have taken the loss function to be known. However, it is clear that loss misspecification is also an important element of robust decision making. Further work is needed to develop a unified approach for dealing with this. Our framework ignores misspecification in the loss function. Certain loss functions are often chosen for computational ease or because they posses other desirable properties such as convexity. Also, elicitation of the true loss function can be difficult (for an example see the application discussed in section 5.1). Hence for completeness, a robustness analysis of a decision system must take this into account. Ruggeri *et al.* (2005) (pages 636-639) provides some further discussion and references.

Acknowledgements

This work was partially supported by an EPRSC grant from the Industrial Doctorate Centre (SABS-IDC) at Oxford University and by Hoffman-La Roche. We are grateful to David Draper and Dimitris Fouskakis for providing us with the data from the RAND study on quality of hospital care. We also thank Florian Lipsmeier, George Nicholson and Tristan Gray-Davies for many helpful discussions and suggestions. Holmes gratefully acknowledges support for this research from the Oxford-Man Institute, EPSRC, the Medical Research Council and the Wellcome Trust.

References

Ahmadi-Javid, A. 2011. An information-theoretic approach to constructing coherent risk measures. *Pages 2125–2127 of: Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*. IEEE.

Ahmadi-Javid, A. 2012. Entropic value-at-risk: A new coherent risk measure. *Journal of Optimization Theory and Applications*, **155**(3), 1105–1123.

Beaumont, M.A, Zhang, W., & Balding, D.J. 2002. Approximate Bayesian computation in population genetics. *Genetics*, **162**(4), 2025–2035.

Belsley, D.A., Kuh, E., & Welsch, R.E. 2005. *Regression diagnostics: Identifying influential data and sources of collinearity*. Vol. 571. John Wiley & Sons.

Berger, J.O. 1984. The robust Bayesian viewpoint (with discussion). *Robustness in Bayesian Statistics (J. Kadane, ed.)*, 63–124.

Berger, J.O. 1985. *Statistical decision theory and Bayesian analysis.* Springer.

Berger, J.O. 1994. An overview of robust Bayesian analysis – with discussion. *Test*, **3**(1), 5–124.

Berger, J.O., & Berliner, L.M. 1986. Robust Bayes and empirical Bayes analysis with ε-contaminated priors. *The Annals of Statistics*, **14**(1), 461–486.

Bernardo, J.M., & Smith, A.F.M. 1994. *Bayesian Theory.* Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons.

Bissiri, P.G., & Walker, S.G. 2012a. Converting information into probability measures via the Kullback–Leibler divergence. *Annals of the Institute of Statistical Mathematics*, **64**, 1139–1160.

Bissiri, P.G., & Walker, S.G. 2012b. On Bayesian learning from loss functions. *Journal of Statistical Planning and Inference*, **142**, 3167–3173.

Bissiri, P.G., Holmes, C.C., & Walker, S.G. 2013. *A General Framework for Updating Belief Distributions.* http://arxiv.org/abs/1306.6430.

Box, George E.P., & Draper, N.R. 1987. *Empirical model-building and response surfaces.* John Wiley & Sons.

Breuer, T., & Csiszár, I. 2013a. Systematic stress tests with entropic plausibility constraints. *Journal of Banking & Finance*, **37**(5), 1552–1559.

Breuer, Thomas, & Csiszár, Imre. 2013b. Measuring Distribution Model Risk. *arXiv preprint arXiv:1301.4832.*

Brown, BJ, Fearn, Tom, & Vannucci, Marina. 1999. The choice of variables in multivariate regression: a non-conjugate Bayesian decision theory approach. *Biometrika*, **86**(3), 635–648.

Chipman, H.A., George, E.I., & McCulloch, R.E. 1998. Bayesian CART model search. *Journal of the American Statistical Association*, **93**(443), 935–948.

Dalalyan, A., & Tsybakov, A.B. 2008. Aggregation by exponential weighting, sharp PAC-Bayesian bounds and sparsity. *Machine Learning*, **72**, 39–61.

Dalalyan, A., & Tsybakov, A.B. 2012. Sparse regression learning by aggregation and Langevin Monte-Carlo. *Journal of Computer and System Sciences*, **78**, 1423–1443.

Dempster, A.P. 1975. A subjectivist look at robustness. *Bull. Internat. Statist. Inst*, **46**, 349–374.

Denison, D., Holmes, C.C., Mallick, B., & A.F.M., Smith. 2002. *Bayesian methods for nonlinear classification and regression.* Wiley.

Draper, David, Kahn, Katherine L, Reinisch, Ellen J, Sherwood, Marjorie J, Carney, Maureen F, Kosecoff, Jacqueline, Keeler, Emmett B, Rogers, William H, Savitt, Harry, Allen, Harris, *et al.* 1990. Studying the Effects of the DRG-Based Prospective Payment System on Quality of Care: Design, Sampling, and Fieldwork. *JAMA: The Journal of the American Medical Association*, **264**(15), 1956–1961.

Fearnhead, P., & Prangle, D. 2012. Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **74**(3), 419–474.

Fouskakis, D., Ntzoufras, I., & Draper, D. 2009. Bayesian Variable Selection Using Cost-Adjusted BIC, with Application to Cost-Effective Measurement of Quality of Health Care. *The Annals of Applied Statistics*, **3**(2), pp. 663–690.

Fouskakis, Dimitris, & Draper, David. 2008. Comparing stochastic optimization methods for variable selection in binary outcome prediction, with application to health policy. *Journal of the American Statistical Association*, **103**(484), 1367–1381.

Gelman, Andrew. 2007. *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press.

Gilboa, I., & Schmeidler, D. 1989. Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, **18**(2), 141–153.

Good, I.J. 1952. Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 107–114.

Hand, David J. 2006. Classifier technology and the illusion of progress. *Statistical Science*, **21**(1), 1–14.

Hansen, Lars Peter, Sargent, Thomas J, Turmuhambetova, Gauhar, & Williams, Noah. 2006. Robust control and model misspecification. *Journal of Economic Theory*, **128**(1), 45–90.

Hansen, L.P., & Sargent, T.J. 2001a. Acknowledging misspecification in macroeconomic theory. *Review of Economic Dynamics*, **4**(3), 519–535.

Hansen, L.P., & Sargent, T.J. 2001b. Robust control and model uncertainty. *The American Economic Review*, **91**(2), 60–66.

Hansen, L.P., & Sargent, T.J. 2008. *Robustness*. Princeton university press.

Hjort, N.L., Holmes, C.C., Muller, P., & Walker, S.G. 2010. *Bayesian nonparametrics*. Cambridge University Press.

Huber, Peter J. 2011. *Robust statistics*. Springer.

Kadane, J.B. (ed). 1984. *Robustness of Bayesian analyses*. Vol. 4. North Holland.

Kadane, J.B., & Chuang, D.T. 1978. Stable decision problems. *The Annals of Statistics*, 1095–1110.

Kadane, J.B., & Srinivasan, C. 1994. Discussion of Berger, J.O., An overview of robust Bayesian analysis – with discussion. *Test*, **3**(1), 116–120.

Kahn, Katherine L, Rubenstein, Lisa V, Draper, David, Kosecoff, Jacqueline, Rogers, William H, Keeler, Emmett B, & Brook, Robert H. 1990. The effects of the DRG-based prospective payment system on quality of care for hospitalized Medicare patients. *Journal of the American Medical Association*, **264**, 1953–1955.

Marin, JM., Pudlo, P., Robert, C.P., & Ryder, R.J. 2012. Approximate Bayesian computational methods. *Statistics and Computing*, **22**(6), 1167–1180.

Marjoram, P., Molitor, J., Plagnol, V., & Tavaré, S. 2003. Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, **100**(26), 15324–15328.

Marmot, M.G., *et al.* 2012. The benefits and harms of breast cancer screening: an independent review. *Lancet*.

Minka, Thomas P. 2001. Expectation propagation for approximate Bayesian inference. *Pages 362–369 of: Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc.

National Research Council:, Committee on the Analysis of Massive Data, Committee on Applied and Theoretical Statistics, Board on Mathematical Sciences and Their Applications, & Division on Engineering and Physical Sciences. 2013. *Frontiers in Massive Data Analysis*. The National Academies Press.

Parmigiani, Giovanni. 1993. On optimal screening ages. *Journal of the American Statistical Association*, **88**(422), 622–628.

Parmigiani, Giovanni, & Inoue, Lurdes Y. T. 2009. *Decision Theory*. John Wiley & Sons, Ltd.

Polson, Nicholas G., Scott, James G., & Windle, Jesse. 2013. *Bayesian inference for logistic models using Polya-Gamma latent variables*. Most recent version: Feb. 2013.

Pritsker, M. 1997. Evaluating value at risk methodologies: accuracy versus computational time. *Journal of Financial Services Research*, **12**(2-3), 201–242.

Rasmussen, C.E., & Williams, C.K.I. 2006. Gaussian processes for machine learning.

Ratmann, O., Andrieu, C., Wiuf, C., & Richardson, S. 2009. Model criticism based on likelihood-free inference, with an application to protein network evolution. *Proceedings of the National Academy of Sciences*, **106**(26), 10576–10581.

Rios Insua, D., & Ruggeri, F. (eds). 2000. *Robust Bayesian Analysis*. Springer.

Robbins, H. 1952. Asymptotically Sub-Minimax Solutions of the Compound Decision Problem'in J. *Page 13 of: Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*.

Robert, C.P, & Casella, G. 2004. *Monte Carlo statistical methods*. Vol. 319. Citeseer.

Rockafellar, R.T, & Uryasev, S. 2000. Optimization of conditional value-at-risk. *Journal of risk*, **2**, 21–42.

Rostek, M. 2010. Quantile Maximization in Decision Theory. *The Review of Economic Studies*, **77**(1), 339–371.

Rue, Håvard, Martino, Sara, & Chopin, Nicolas. 2009. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the royal statistical society: Series b (statistical methodology)*, **71**(2), 319–392.

Ruggeri, F., & Wasserman, L. 1993. Infinitesimal sensitivity of posterior distributions. *Canadian Journal of Statistics*, **21**(2), 195–203.

Ruggeri, F., Ríos Insua, D., & Martín, J. 2005. Robust Bayesian Analysis. *Handbook of statistics*, **25**, 623–667.

Savage, L.J. 1954. *The foundations of statistics*. New York: Wiley.

Shapiro, S, Venet, W, Strax, P, & Venet, L. 1988. Periodic screening for breast cancer: the Health Insurance Plan project and its sequelae, 1963-1986. *Baltimore, Maryland: The John Hopkins University Press*.

Sivaganesan, S. 1994. Discussion of Berger, J.O., An overview of robust Bayesian analysis – with discussion. *Test*, **3**(1), 116–120.

Sivaganesan, S. 2000. Global and local robustness approaches: uses and limitations. *Pages 89–108 of:* Rios Insua, D., & Ruggeri, F. (eds), *Robust Bayesian Analysis*. Springer.

Varin, Cristiano, Reid, Nancy, & Firth, David. 2011. An overview of composite likelihood methods. *Statistica Sinica*, **21**(1), 5–42.

Vidakovic, B. 2000. Γ-minimax: a paradigm for conservative robust Bayesians. *Pages 241–259 of:* Rios Insua, D., & Ruggeri, F. (eds), *Robust bayesian analysis*. Springer.

Von Neumann, J., & Morgenstern, O. 1947. *The theory of games and economic behavior*. Princeton university press.

Wainwright, M., & Jordan, M.I. 2003. 2003. Graphical models, exponential families and variational inference. *Foundations and Trends in Machine Learning*, 1–305.

Wald, A. 1950. *Statistical decision functions.* Wiley.

Wasserman, L. 1992. Recent methodological advances in robust Bayesian inference. **4**, 483–502.

Whittle, P. 1990. *Risk-sensitive Optimal Control*. Wiley.

Wu, Dongfeng, Rosner, Gary L, & Broemeling, Lyle D. 2007. Bayesian inference for the lead time in periodic cancer screening. *Biometrics*, **63**(3), 873–880.

Zhang, T. 2006a. From ϵ-entropy to KL-entropy: Analysis of minimum information complexity density estimation. *Annals of Statistics*, **34**, 2180–2210.

Zhang, T. 2006b. Information theoretical upper and lower bounds for statistical estimation. *IEEE Trans. Inform. Theory*, **52**, 1307–1321.

A Proof of Theorem 4.2 in Section 4.

Reproduced and amended from (Bissiri *et al.*, 2013).

Assume that Θ contains at least two distinct points, say θ_1 and θ_2. Otherwise, π is degenerate and the thesis is trivially satisfied. To prove this theorem, it is sufficient to consider the case $n = 2$ and a very specific choice for π, taking $\pi = p_0\delta_{\theta_1} + (1 - p_0)\delta_{\theta_2}$, where $0 < p_0 < 1$. Any probability measure absolutely continuous with respect to π has to be equal to $p\delta_{\theta_1} + (1 - p)\delta_{\theta_2}$, for some $0 \leq p \leq 1$. Therefore, in this specific situation, the cost function, $l(\cdot) = \left\{E_\pi[-L(\theta)] + \lambda^{-1}g(\pi \parallel \pi_I)\right\}$, to be minimised becomes:

$$\begin{aligned} l(p, p_0, L_I) := {} & p\, L_I(\theta_1) \,+\, (1-p)\, L_I(\theta_2) \\ & +\, p\, g\left(\frac{p}{p_0}\right) \,+\, (1-p)\, g\left(\frac{1-p}{1-p_0}\right), \end{aligned}$$

where g is a divergence measure, $L_I(\theta_i) = L(\theta_i, I_1) + L(\theta_i, I_2)$ for data $I = (I_1, I_2)$ and $L_I(\theta_i) = L_1(\theta_i, I_j)$ for $I = I_j$, $i, j = 1, 2$. Denote by p_1 the probability $\pi_{I_1}(\{\theta_1\})$, i.e. the minimum point of $l(p, p_1, L_{(I_1,I_2)})$ as a function of p, and by p_2 the probability $\pi_{(I_1,I_2))}(\{\theta_1\})$. By hypotheses, p_2 is the unique minimum point of both loss functions $l(p, p_1, L_{I_2})$ and $l(p, p_0, L_{(I_1,I_2)})$. Again by hypothesis, we shall consider only those functions L_{I_1} and L_{I_2} such that each one of the functions $l(p, p_0, L_{I_1})$, $l(p, p_1, L_{I_2})$, and $l(p, p_0, L_{(I_1,I_2)})$, as a function of p, has a unique minimum point, which is p_1 for the first one and p_2 for the second and third one. The values p_1 and p_2 have to be strictly bigger than zero and strictly smaller than one: this was proved by Bissiri and Walker (2012) in their Lemma 2. Hence, p_1 has to be a stationary point of $l(p, p_0, h_{I_1})$ and p_2 of both the functions $l(p, p_1, L_{I_2})$ and $l(p, p_0, L_{(I_1,I_2)})$. Therefore,

$$g'\left(\frac{p_1}{p_0}\right) - g'\left(\frac{1-p_1}{1-p_0}\right) = L_{I_1}(\theta_2) - L_{I_1}(\theta_1), \tag{14}$$

$$g'\left(\frac{p_2}{p_0}\right) - g'\left(\frac{1-p_2}{1-p_0}\right) = L_{(I_1,I_2)}(\theta_2) - L_{(I_1,I_2)}(\theta_1), \tag{15}$$

$$g'\left(\frac{p_2}{p_1}\right) - g'\left(\frac{1-p_2}{1-p_1}\right) = L_{I_2}(\theta_2) - L_{I_2}(\theta_1). \tag{16}$$

Recall that $L_{(I_1,I_2)} = L_{I_2} + L_{I_1}$. Therefore, summing up term by term (14) and (16), and considering (15), one obtains:

$$\begin{aligned} & g'\left(\frac{p_2}{p_0}\right) - g'\left(\frac{1-p_2}{1-p_0}\right) \\ & \qquad = g'\left(\frac{p_1}{p_0}\right) - g'\left(\frac{1-p_1}{1-p_0}\right) + g'\left(\frac{p_2}{p_1}\right) - g'\left(\frac{1-p_2}{1-p_1}\right). \end{aligned} \tag{17}$$

Recall that by hypothesis (14)–(16) need to hold for every two functions L_{I_1} and L_{I_2} arbitrarily chosen with the only requirement that p_1 and p_2 uniquely exist. Hence, (17) needs to hold for every (p_0, p_1, p_2) in $(0, 1)^3$. By substituting $t = p_0$, $x = p_1/p_0$ and $y = p_2/p_1$, (17) becomes

$$\begin{aligned} & g'(xy) - g'\left(\frac{1-txy}{1-t}\right) \\ & \qquad = g'(x) - g'\left(\frac{1-tx}{1-t}\right) + g'(y) - g'\left(\frac{1-txy}{1-tx}\right), \end{aligned} \tag{18}$$

which holds for every $0 < t < 1$, and every $x, y > 0$ such that $x < 1/t$ and $y < 1/(xt)$. Being g convex and differentiable, its derivative g' is continuous. Therefore, letting t go to zero, (18)

implies that

$$g'(xy) = g'(x) + g'(y) - g'(1) \tag{19}$$

holds true for every $x, y > 0$. Define the function $\varphi(\cdot) = g'(\cdot) - g'(1)$. This function is continuous, being g' such, and by (19), $\varphi(xy) = \varphi(x) + \varphi(y)$ holds for every $x, y > 0$. Hence, $\varphi(\cdot)$ is $k \ln(\cdot)$ for some k, and therefore

$$g'(x) = k \ln(x) + g'(1), \tag{20}$$

where $k = (g'(2) - g'(1))/\ln(2)$. Being g convex, g' is not decreasing and therefore $k \geq 0$. If $k = 0$, then g' is constant, which is impossible, otherwise, for any h_I, p_1 satisfying (14) either would not exist or would not be unique. Therefore, k must be positive. Being $g(1) = 0$ by assumption, (20) implies that $g(x) = k\,x \ln(x) + (g'(1) - k)(x - 1)$. Hence,

$$g(\pi_1, \pi_2) = k \int \ln\left(\frac{\mathrm{d}\pi_1}{\mathrm{d}\pi_2}\right) \mathrm{d}\pi_1$$

holds true for some $k > 0$ and for every couple of measures (π_1, π_2) on Θ such that π_1 is absolutely continuous with respect to π_2.

B Actions Selected by Simulated Annealing Search

Table 1: List of variables selected by our method. The first two columns give the index number and name of the variables selected by our simulated annealing algorithm. Column 3 indicates the fixed cost of inclusion of each variable, units in minutes.

Variable			Method					
			Action 1	Action 2	Action 3	Action 4	Action5	Action 6
Index	Name	Cost						
1	Systolic BP score	0.5		*	*	*	*	
2	Age	0.5			*	*	*	*
3	Blood urea nitrogen	1.5						*
4	APACHE II coma score	2.5	*	*				*
12	Initial temperature	0.5	*	*	*			*
14	Chest pain day 1?	0.5					*	*
35	Hepatobiliary history?	0.5	*	*	*	*	*	
36	Renal history score	1.0					*	
37	APACHE respiratory rate score	1.0	*		*	*		*
51	Confusion day 1	0.5	*	*				*
54	Pulmonary edema score	0.5						*
61	Wheezing at admission?	0.5				*		
62	Blood system count	2.5	*	*	*	*	*	
68	Co-morbid smoking score	0.5	*					
70	APACHE pH score	1.0					*	
74	Cardiac history score	0.5	*	*	*	*	*	*
75	Neurologic history score	0.5	*	*	*	*	*	*
76	Oncologic history score	0.5	*	*	*	*	*	
77	Immunologic history score	0.5			*	*	*	
78	Musculoskeletal score	0.5	*	*	*	*		